



AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search

<http://ageconsearch.umn.edu>

aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

Ranking Environmental Projects

Please see revised/updated version of this paper - Working Paper 1507 published April 2015 at <http://purl.umn.edu/204305>.

David J. Pannell

Centre for Environmental Economics and Policy
School of Agricultural and Resource Economics,
The University of Western Australia,
Crawley, WA 6009, Australia

*E-mail address: david.pannell@uwa.edu.au

5th September 2013
Working Paper 1312
School of Agricultural and Resource Economics
<http://www.are.uwa.edu.au>



THE UNIVERSITY OF
WESTERN AUSTRALIA
Achieve International Excellence

Citation: Pannell, D.J., (2013) *Ranking Environmental Projects* Working Paper 1312, School of Agricultural and Resource Economics, University of Western Australia, Crawley, Australia.

© Copyright remains with the author of this document.

Ranking Environmental Projects

Contents

1. What is being ranked?	4
2. Benefits/Costs	5
3. Benefits	7
3.1 With versus without	7
3.2 Environmental condition and values	9
3.3 Estimating and measuring values	11
3.4 Adoption and compliance	15
3.5 Project risks	18
3.6 Time and time lags	20
3.7. Multiple benefits	24
3.8 Private benefits	26
3.9 Scoring variables	28
4. Costs	30
4.1 Cash costs	30
4.2 In-kind costs	31
4.3 Private costs	32
4.4 Maintenance costs	34
4.5 Sharing costs	35
5. Uncertainty	36
6. Simplifications	38
7. Mistakes to avoid	39
8. Summary	42
References	45

Abstract

Environmental agencies and organisations face the challenge of deciding which of the many possible environmental projects they should support with their limited resources. Projects vary greatly in environmental benefits and costs, so selecting the best projects can make a major difference to the level of environmental benefits that can be generated for a given budget. Key principles for ranking environmental projects are presented and explained. A suitable formula to use as a metric for ranking projects is developed and explained. The formula accounts for environmental values, the effectiveness of management, time lags, behaviour change, various risks and various costs. The suggested formula includes a number of appropriate simplifications. A number of common mistakes to avoid are outlined.

Key words: environment, economics, project prioritisation, uncertainty, behaviour change, risk, environmental values, technical feasibility

JEL Codes: D82, Q20, Q28

The funding available for environmental projects and policies is usually a small percentage of the money we would need to deal comprehensively with all environmental problems. As a result, whether we like it or not, we have to choose what we do and don't protect. Even programs that don't explicitly prioritise their environmental investments do so implicitly – they just do it in a non-transparent, and usually very poor, way.

In my experience, the difference in potential environmental outcomes between poor prioritisation processes and good ones is enormous.

Doing a good job of ranking the investment options is not that hard if you are aware of a few principles, but it seems to me that most people who are responsible for deciding how environmental funds get allocated are not aware of these principles. Indeed, some of the most commonly used approaches to ranking environmental projects are guaranteed to result in very poor rankings. As a result, we miss easy opportunities to deliver much greater environmental outcomes.

My aim in this document is to outline a set of relevant principles and insights that will help environmental decision makers choose the best projects. My focus is on collecting and analysing the information needed to provide high-quality project rankings. There is another set of issues about how the rules of the program are designed to provide incentives for its participants to behave appropriately (e.g. Pannell and Roberts 2010), but I won't be covering those here. I'll be talking about information, calculations and clear thinking in the process of choosing which projects to support.

My aim is to help with practical decision making. As a result, I'll be talking about the possibility of cutting corners by simplifying aspects of the process. You'll see that I'm not averse to well-considered simplifications, but very wary of the risk that some simplifications will sabotage the

whole process. For a practical system, simplifications are essential, but bad simplifications are disastrous.

Throughout, I will be assuming that the aim is to provide the most valuable environmental outcomes for the available resources.

1. What is being ranked?

The first requirement is to be clear about what is being ranked. Sometimes programs set out to rank a set of projects that they might invest in. The projects should define what would be done, to which environmental assets, where, and by whom. (I'll use the term "environmental asset" to refer to any identifiable feature, entity, place, or species that might become a target for investment.)

At other times, programs seek to rank a set of environmental assets, with no explicit project activities defined. There is a risk here – if you don't define the project activities for an environmental asset, you cannot rank them on the basis of providing the most valuable outcomes.

The problem is that the environmental value for money depends on the answers to questions like, "what is the technical feasibility of protecting the asset?", "to what extent would the community cooperate?" and "what would it cost to protect the asset?" However, those questions can only be answered for a particular set of actions or interventions.

To further illustrate the point, various different projects could be defined for the same environmental asset. One potential project might have very ambitious goals, aiming to return the asset to pristine condition, while another might aim for a moderate improvement in its condition. Some of these different projects for the same asset may offer relatively good value for money while others don't (e.g. Roberts et al. 2012). So you cannot conclude that investing in any particular asset is good or bad without being clear about the project actions that will be undertaken.

If the analysis is limited to environmental assets, not projects, then it is important to be aware of what can and cannot be done with the results. What you can reasonably do is filter the assets to identify ones where it is relatively likely that a well-designed project would deliver worthwhile benefits. This could be done using variables such as:

- the value or significance of the assets,
- the levels of degradation they have already suffered or are likely to suffer in future, and
- the feasibility of managing them (in a loose general sense that doesn't require specification of particular management actions).

You should not be making final decisions about which assets received funding, because that does require the specification of projects. Rather, you would be concluding that some assets are probably not worth considering further, and so not worth developing projects for.

Even this is not without risks. Because you are not looking at all of the relevant information, there is a chance of excluding some assets that would actually be worth investing in. For example, you might exclude investment in a particular asset because it seems likely to provide only modest benefits, but

if the cost of the project is low enough, it could still be worth doing. With this process of filtering assets, you would miss out on cases like that.

However, it still might be worth filtering assets as part of a more comprehensive process. Indeed that is exactly what we do in Step 2 of INFFER (the Investment Framework for Environmental Resources) (Pannell et al. 2012). This is a simplification we use to reduce the cost of the system. If we can knock out some potential investments based on partial information, it takes less work to properly evaluate and rank a reduced set of potential projects.

If you *must* make final investment decisions based on assets, not projects, you need to imagine a notional project for each asset. Even a rough-and-ready notional project definition would be better than nothing.

2. Benefits/Costs

Suppose you manage an environmental program that has a budget available for spending on environmental projects and there is not enough money to fund every proposed project. You have to decide which projects to fund. How should you do it?

The first principle is that projects should be ranked using a metric (a formula) that consists of a measure of project benefits divided by a measure of project costs. Economists call this metric a Benefit: Cost Ratio (BCR).

$$BCR = B/C$$

There are plenty of project ranking metrics out there in actual use that don't do this. Some subtract costs instead of dividing them, and some (remarkably) ignore costs entirely. These are mistakes that are costly to the environment.

To illustrate, consider the following three hypothetical projects, with the indicated benefits (B) and costs (C). Because the budget is limited, the first project we should choose is the one with the highest benefits per unit cost (the highest BCR) = project 1. But if we rank according to $B - C$ the top ranked project seems to be project 2, while ranking according to B (ignoring costs) tells us that project 3 is best.

Project	B	C	BCR	$B - C$	Rank(BCR)	Rank($B - C$)	Rank(B)
1	5	1	5	4	1	2	3
2	7	2	3.5	5	2	1	2
3	8	7	1.1	1	3	3	1

The loss of environmental values from using the wrong metric (i.e., ranking according to $B - C$ or B) depends on how tight the budget is. Assuming that the budget is enough to fund 10% of projects, the loss of environmental benefits is 12% for $B - C$ and 19% for B (based on simulating 1000 funding rounds with 100 potential projects in each).

In other words, fixing up the formula is like increasing the program budget by 14% or 23%. It's much easier to fix the formula than to increase the budget!

In the examples above, I've assumed that we know what the benefits and costs would be for each project. Later sections in this paper will deal in detail with how we should estimate the benefits and costs. For now I'll just make these two observations.

The benefits used in the ranking metric should be the benefits of the proposed intervention or project, not the total benefits of the environmental asset. What difference can be made by the intervention or project, and how important is that difference?

The costs should also represent the costs of the intervention or project. If this project did not go ahead, what level of resources could be diverted to other uses?

Sometimes people criticise the use of the Benefit: Cost Ratio (BCR) to rank projects, on the basis that that it can be manipulated to some extent by moving costs between the denominator and the numerator (e.g. Office of Best Practice Regulation, 2009; Jenkins et al., 2011). For example, suppose you have already calculated an initial BCR for a project, but now you find that there is an additional cost that should be included. You could do one of two things with that cost: you could subtract it from the numerator, resulting in smaller benefits in the BCR, or you could add it to the denominator, resulting in larger costs in the BCR. If benefits exceed costs even after accounting for the new cost, then subtracting the new cost from the numerator would result in a larger BCR than adding it to the denominator.

For example, suppose that a proposed project has benefits of \$10m, program costs of \$2m (requested from the funding program) and other costs of \$1m (from other sources, such as the private sector). We could potentially calculate the BCR as $10/(2+1) = 3.3$, or else as $(10-1)/2 = 4.5$.

This criticism of using the BCR for ranking projects reveals a lack of understanding of the logic of the formula. To rank projects correctly, the costs that go in the denominator are the costs that would be drawn from a limited pool of funds. Any costs that are not drawn from a limited pool should, in principle, be subtracted from the numerator, rather than being added to the denominator. It is not correct to move costs arbitrarily between the two. There is a clear logic about which costs go where. It's surprising how often this misconception is repeated, even by economists. So if the other costs are not drawn from a fixed budget, the correct procedure is to subtract the other costs from the benefits, meaning that the correct BCR for the above project would be 4.5.

Things get a bit tricky, however, if projects also require ongoing maintenance funding beyond the current project, and the budget for maintenance funding is expected to be fully allocated. This is realistic for many (probably most) projects. In this case, there are actually two constraints that must be satisfied: the current program budget and the long-term maintenance budget. Strictly, in this situation, projects cannot be ranked using a single formula as a metric. The program would need a mathematical programming model to select which projects deliver the most benefits while satisfying both constraints. In practice, after testing various approaches, I believe that a reasonable approximation is to add up both costs (short-term program costs and long-term maintenance costs) and include the total as the denominator in the single formula.

3. Benefits

In this section we will cover a number of points about the estimation of benefits from an environmental project. Initially, to keep things simple, I'll talk about the case where there is a single type of benefit being generated by an environmental project (e.g. a threatened species is being made safer). In later sections I'll talk about cases with multiple types of benefits from the same project.

3.1 With versus without

This first point is deceptively simple. It is that the benefit of an environmental project is the change in value of the environmental asset as a result of the project. In other words, it is a difference: the difference between the environmental value with the project and without the project.

So, to estimate the benefits of a project, you need two pieces of information: the environmental values with the project and the values without the project. Usually, when we are evaluating a project, the project has not yet been implemented. In that case, both of the required pieces of information have to be predicted. You can't observe them, because they are in the future.

Note that comparing environmental values "with versus without" the project is not the same as comparing values "before versus after" the project. The reason is that the condition of the asset would probably not be static in the absence of the project. For example, it may be that the asset would degrade in the absence of the project, but its condition would be improved by the project (relative to its current condition). This is illustrated in Figure 1.

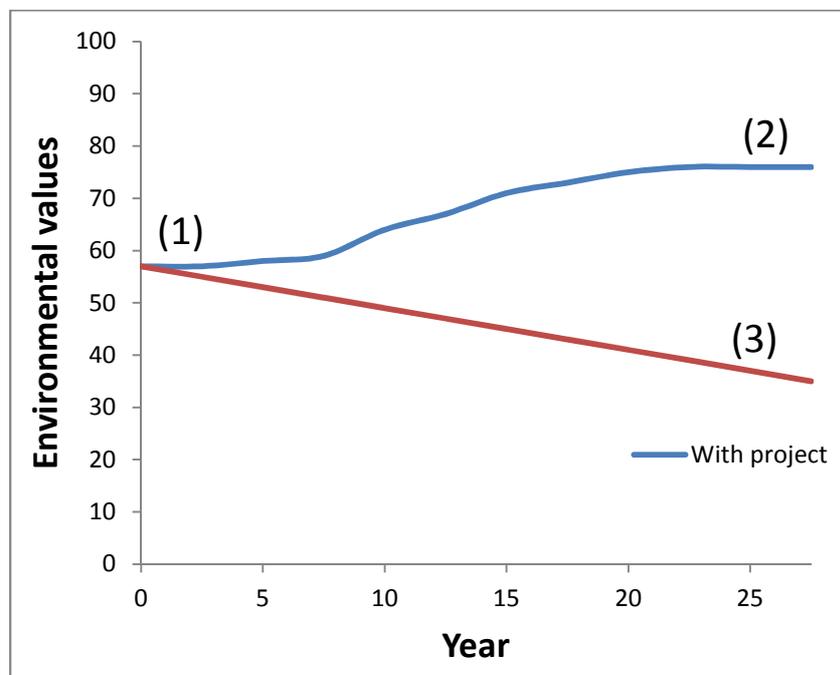


Figure 1.

The graph illustrates a case where the asset currently has a value of 57 [labelled (1)]. (The 57 is just some measure of value – we’ll discuss values more in later sections.) Without the proposed project, the value is expected to decline steadily, to a score of 37 after 25 years [labelled (3)]. With the project, value would increase to a score of 76 after 25 years [labelled (2)].

Clearly, in this example, the benefits of the project grow over time (the two lines diverge in Figure 1). Ideally, we would estimate the benefits in each year after the project is implemented and add them up (after allowing for discounting, which we’ll cover in a later section). A practical simplification is to estimate the environmental benefits based on the difference in the asset value with and without the project in a particular future year. For example, we might choose to focus on 25 years in the future, and estimate values at that date with and without the project. In doing this, we need to be careful that we deal appropriately with time (see a later section for details).

Assuming we go with that simplified approach (focusing on benefits at year 25), the relevant measure of project benefits for ranking projects is (2) minus (3). I have seen ranking systems which use (1) minus (3), (2) minus (1), (1) alone or (2) alone, and sometimes more than one of these in the same ranking system, but they are all irrelevant. If you include (2) minus (3) you should not include any of the others listed. To do so will just make the rankings worse.

Because of the “with versus without” principle, a project can generate benefits even if it does not completely prevent degradation of the environmental asset. As long as it slows or reduces degradation, this should be measured as a benefit. Figure 2 shows an illustration of this. In this example, future asset condition with the project (2) is below the initial asset condition (1), but is above future asset condition without the project (3). Since the project benefit is (2) minus (3), the benefit is positive.

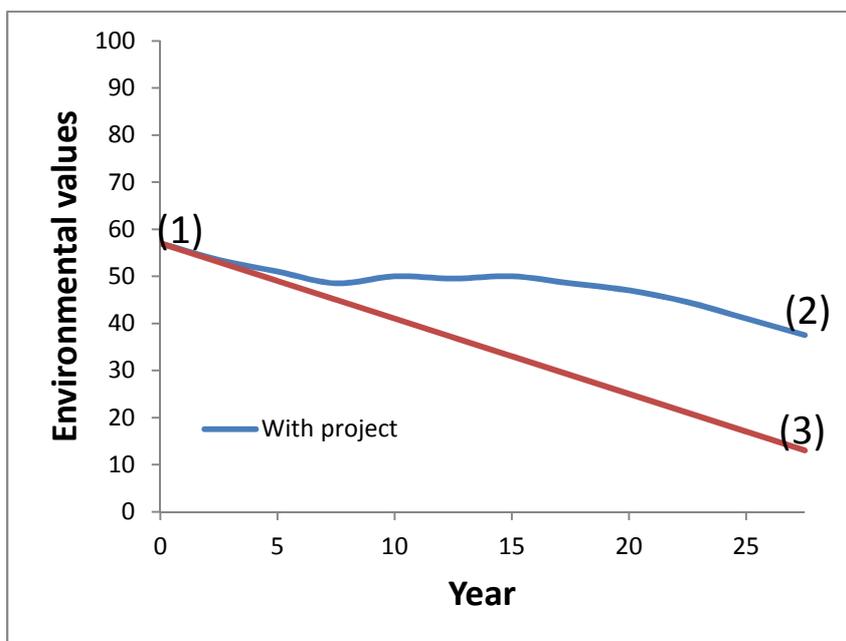


Figure 2.

On the other hand, a project that superficially appears to generate large benefits may actually not do so, because those benefits would have been generated even without the project. In other words, the benefits are not 'additional' to what would have happened anyway. The without-project line in the graph would be almost as high as the with-project line, so the difference between them (= the benefit of the project) would be minimal (Figure 3).

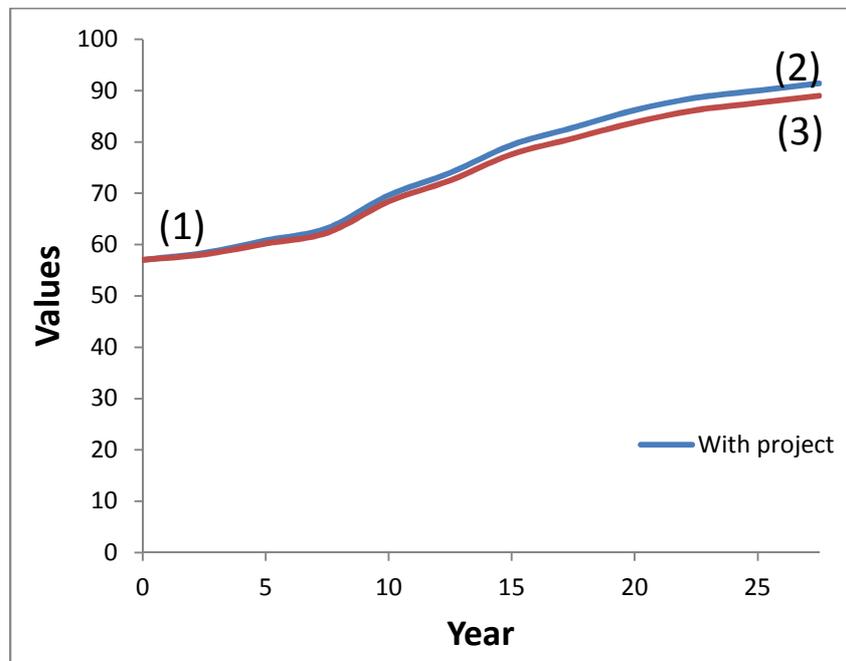


Figure 3.

For example, suppose that a proposed project encourages farmers to adopt a new type of environmentally beneficial crop, where that crop is much more profitable than farmers' existing crops. If the private benefits are large enough, it's a safe bet that the farmers would have adopted the new crop even without the project. It would have been promoted by word of mouth and by private farm business consultants. Adoption of the crop for commercial reasons would have generated environmental benefits as a spin-off.

Making good predictions about the "without project" scenario can be quite difficult, requiring good knowledge of the environment, the relevant management practices and the people whose behaviour matters. Weak thinking about the "without" scenario for environmental projects is a common failing, sometimes leading to exaggerated estimates of the benefits.

3.2 Environmental condition and values

In the previous section I said that the benefit from an environmental project is the difference between the environmental values with the project and without the project. This time I will break that down a bit. The point of this section is that there are two parts to that change in environmental

values: a change in the physical condition of the environment, and a resulting change in the values generated by the environment (in other words, the value of the environmental services).

So, to estimate the benefits of an environmental project, you need to (a) predict the physical environmental conditions with and without the project, and (b) translate the difference into a measure of value or importance or significance.

This raises the question, what is the relationship between the physical condition of an environmental asset and the values it provides to the community? As environmental conditions improved, values would increase, but is it a simple linear increase, or something else? To some extent, this would depend on how you measure the condition of the environment, but a common result in the environmental economics literature is for values to increase at a decreasing rate, as illustrated in Figure 4. (In fact, we see this sort of relationship for all sorts of things, not just environmental goods.)

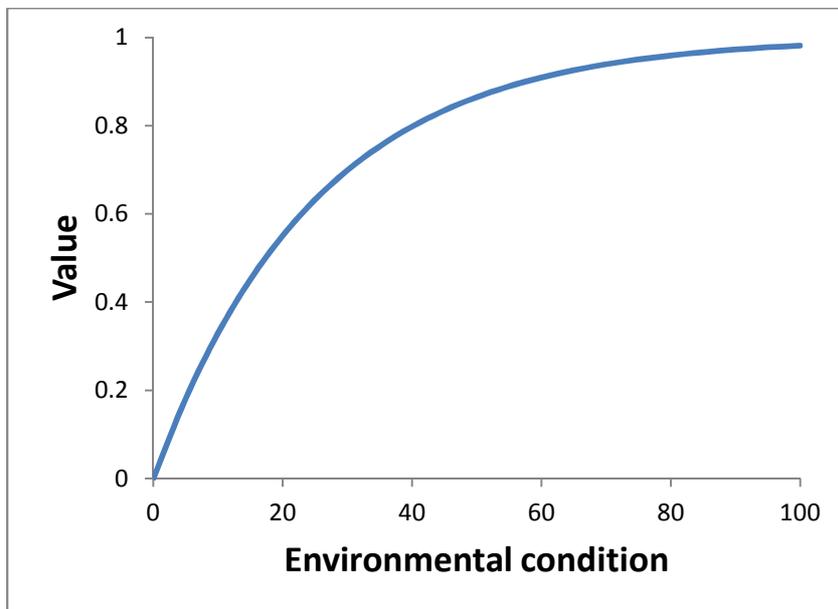


Figure 4.

The relationship in Figure 4 relates to the values that people in the community put on environmental improvements, but it is also consistent with one aspect of the ways that environmental scientists tend to think about values: if something is rarer, each unit of it is considered more valuable.

In theory, if you could quantify environmental conditions and knew the relationship between condition and value, you could read off the change in value from this graph. For example, Figure 5 shows that a project that increases the environmental condition score from 40 to 60 results in an increase in value from about 0.8 to 0.9. If we are measuring the value in millions of dollars, the benefit of that project would be \$100,000.

$$\text{Benefit} = V(P_1) - V(P_0) = V(60) - V(40) = 0.9 - 0.8 = 0.1 \text{ \$million} = \$100,000$$

where V is value, which depends on the physical condition, P_1 is physical condition with the project and P_0 is physical condition without the project.

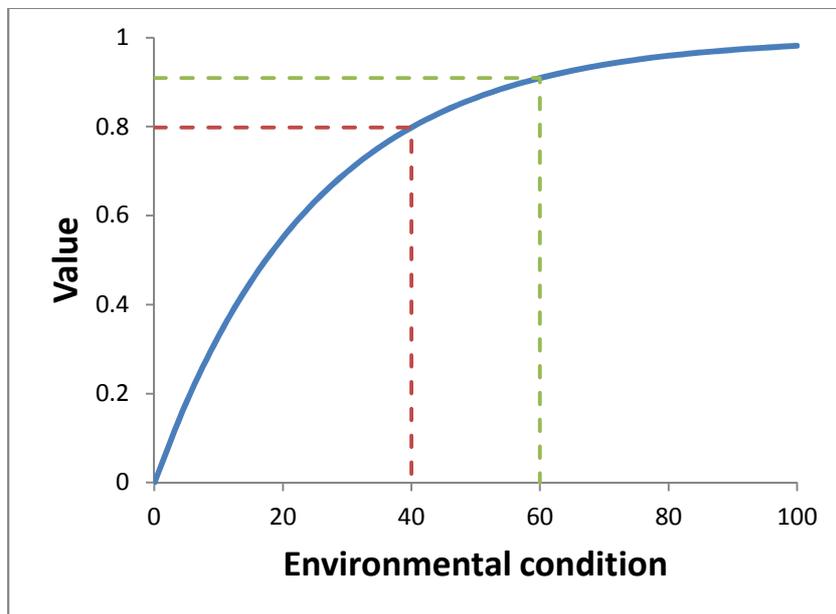


Figure 5.

In practice, we may or may not have a system for quantitatively scoring the type of environmental condition we are interested in in a particular case, but we should at least be able to describe the environmental conditions in words, with and without the project. Then we have to translate those into a measure of value (the topic of the next section).

Amongst the systems I've seen in use for ranking environmental projects, a surprising number make no attempt to evaluate the difference that the project will make to environmental conditions. Without that, there is no prospect of obtaining a meaningful estimate of the benefits from the project, so decision making (and ultimately the environment) suffers.

3.3 Estimating and measuring values

We've seen that measuring the benefits of an environmental project requires attention to two aspects: the change in the physical condition of the environment, and the resulting change in the values generated by the environment (Section 3.2). Suppose we have information about the change in physical conditions. How should we convert that to a measure of value or importance that we can use to rank projects? We need to do this in a way that is consistent between the different projects that we'll want to compare.

Let's consider three options, which are quite different in nature, but which are all actually used in real-world environmental programs.

(a) Scientific principles

Scientists often use rules of thumb (sometimes captured in an 'Environmental Benefits Index') to evaluate the relative importance of different potential environmental investments. An Australian example is the 'habitat hectares' concept, which is used by the state government in Victoria to evaluate proposed projects. A US example is the Environmental Benefits Index developed by the Natural Resources Research Institute (NRRI) of the University of Minnesota Duluth (<http://beaver.nrri.umn.edu/EcolRank/>). This consists of measures of soil quality risk, water quality risk and habitat quality, each scored out of 100, and then added up to give a total score out of 300.

Key strengths of this approach include:

- The index is based on relatively sound knowledge of the natural systems.
- Once the system has been developed, the approach is relatively efficient to apply to many potential projects.

But it also has some weaknesses:

- The resulting Index scores reflect the values of experts, and there is plenty of evidence that experts and the general community sometimes think differently about what is important.
- Environmental Benefits Indexes are set up to evaluate particular types of environmental benefits and cannot evaluate projects that generate different types of benefits. For example, the NRRI's Index is no use for evaluating projects that protect threatened species or reduce air pollution. They can only rank projects of a reasonably similar type.
- Often Environmental Benefits Indexes are not designed in a way that allows the required with-versus-without the project comparison. The NRRI index is an example. Even if we know what difference the project will make to environmental condition, this index would not help us value that difference. This could potentially be addressed by improving the design of the Index, although that would require considerable effort and resources.
- Any system based on scoring, rather than dollars, cannot tell us whether the benefits of a project would exceed its costs. It can tell us how projects should be ranked, but not where the cut-off line should be for projects that are or are not worth funding. In most cases where projects are being ranked, this is not a serious problem because the overall budget is already determined. From a practical perspective, the relevant cut-off line is where the money runs out.

(b) Deliberative processes

A "deliberative process" is 'a process allowing a group of actors to receive and exchange information, to critically examine an issue, and to come to an agreement which will inform decision making' (Gauvin 2009). It involves discussion, debate, and consideration of all information that is considered relevant. Multi-Criteria Analysis often employs this approach, although other approaches can use it as well.

Strengths:

- There is scope to involve both experts and community members to ensure that both perspectives are considered.

- The approach may be seen by stakeholders as being more transparent than the other approaches
- There is an opportunity for participating non-experts to receive detailed information and to participate in discussion and debate about the issues. This means that the outputs are likely to be better informed and better considered than is possible in survey-based approaches.
- The approach is very flexible. All types of benefits and costs can be considered.
- It is possible to generate a large number of valuations relatively efficiently – certainly more cheaply than conducting non-market valuation surveys for each project.

Weaknesses:

- Participants may have vested interests or particular perspectives and may not reflect broader community interests or concerns.
- While the flexibility of the approach is an advantage up to a point, the lack of theoretical rigour can be a problem, resulting in project rankings that don't actually reflect the participants' own values. In other words, too much flexibility can be a problem, particularly if the process goes beyond just looking at values. For example, when it comes to ranking projects, participants should not be free to choose to include costs in any way other than by dividing them into benefits (see Section 2). Some things that people often choose to do in this space are just wrong (which is why I'm writing this document).
- If the output is a score, rather than a dollar value, the approach cannot tell us whether the benefits of a project would exceed its costs.

(c) Dollar values

Environmental economists put a lot of effort into valuing environmental benefits in dollar terms, using a variety of techniques. (See Pannell Discussions 218 to 221 for details.)

Strengths:

- Of the three approaches, this one is likely to best reflect broad community attitudes. It is more independent and less at risk of reflecting the preferences of vested interest groups.
- It allows comparisons across completely different types of benefits.
- It is more rigorous – less *ad hoc* than scoring-based approaches.
- It allows us to determine whether the benefits of a project outweigh its costs.

Weaknesses:

- Respondents to non-market valuation studies may know very little about the things they are being asked to value.
- Conducting separate valuation studies for each project would be prohibitively expensive. Transferring benefit estimates from other similar projects can help to overcome this problem.
- The survey-based methods have been criticised by some economists for relying on hypothetical questions and for giving results that don't seem plausible in some cases. While this debate is interesting, in practice the quality of information from these surveys is probably higher than some other information we need to include in the process. For

example, information about the cause-and-effect relationship between management and environmental conditions is often very weak indeed.

Which is best?

Some people are quite definite in their preferences for one or another of these approaches, or particularly dislike one of them. In my view, it's not a clear-cut decision. They each have pros and cons, and one's choice of which to use may vary depending on the circumstances. The weaknesses that concern me most are: the inability of many Environmental Benefits Indexes to compare outcomes with and without the project; the excessive flexibility of some deliberative approaches, giving participants the flexibility to do dumb things; and the expense of doing comprehensive valuation surveys.

My advice is to weigh up the pros and cons and use whichever approach makes most sense for a particular program. My caution would be that this advice applies specifically to the part of the process that estimates values. For the other parts of the process, and for decisions about how to combine the various bits of information to inform decisions, see the other sections in this document.

A practical compromise

In developing INFFER (Pannell et al. 2012) we attempted to create an approach to valuation that draws on the combined strengths of the three approaches outlined above, while limiting their weaknesses. The approach we developed:

- Can use scientific information if it is available
- Recognises that the relevant benefit is a difference (with minus without the project)
- Can be elicited in a deliberative process involving both experts and community members
- Can be cheap and quick enough to be practical in cases where there are limited resources for estimating values, or where many valuations are needed in a short time
- Provides dollar values
- Can use results from non-market valuation surveys if available

Here is how it works. Define P' as the physical condition of the environmental asset in good condition. For example, it could be an environmental condition of 100 in Figures 4 and 5 (Section 3.2).

Now $V(P')$ is the value of the environmental asset at condition P' . It includes all the different types of values (market and non-market) that are relevant to this environmental asset. In Figures 4 and 5, if $P' = 100$, $V(P')$ would be \$1 million.

Finally, define W as the difference in values between P_1 (physical condition with the project) and P_0 (physical condition without the project) as a proportion of $V(P')$.

$$W = \frac{V(P_1) - V(P_0)}{V(P')}$$

Then we measure the project benefit as $V(P') \times W$:

$$\textit{Benefit} = V(P') \times W$$

$$\begin{aligned}
&= V(P') \times \frac{V(P_1) - V(P_0)}{V(P')} \\
&= V(P_1) - V(P_0)
\end{aligned}$$

So $V(P') \times W$ is equivalent to the correct measure of benefits, $V(P_1) - V(P_0)$ (as outlined in Section 3.2).

The benefit of re-organising the benefits into $V(P')$ and W is that, in my experience, it helps people think clearly and ask the right questions in a situation where they are not going to conduct a non-market valuation survey. $V(P')$ sets an upper bound for the benefits of the project – obviously, the value of the project can't be more than the value of the environmental asset in good condition.

In INFFER, we ask users to score $V(P')$ relative to a set of examples – a table of well-known environmental assets with suggested $V(P')$ scores. We define $V(P')$ as being worth \$20 million per point. This is often done in a group discussion environment, involving a variety of stakeholders.

A risk with this (and other deliberative processes) is that people may provide values that are too high (e.g. see Pannell Discussion 213). A process of reviewing assumptions and comparing them across projects is needed to reduce this risk.

Defining W as a proportion of $V(P')$ helps to highlight that the benefits of the project must be proportional to the effectiveness of the project, which is often missed when people develop their metric for ranking projects.

For example, suppose there are two alternative projects for Asset A. Project (i) would increase the asset value by a factor of 0.3 and Project (ii) would increase it by 0.6. If everything else is equal, Project (ii) would generate benefits that are twice as large as those from Project (i). The metric has to reflect that. This is achieved by multiplying by W .

Finally, a mistake I've seen is to exclude any measure of values from the ranking process. One senior bureaucrat told me that she was opposed to including them because of the risk of them generating controversy. At other times, people seem to simply overlook them. The consequence of this is that the organisation will tend to bias its funding towards less valuable projects. There is an increased risk that they will incorrectly rank projects addressing less-valuable assets relative to more-valuable assets.

3.4 Adoption and compliance

In Sections 3.1, 3.2 and 3.3 I talked about estimating the benefits of environmental projects, as part of the process of ranking projects. To keep things simple, I focused on the predicted environmental changes and their values, but there are other benefit-related factors that we need to account for too. The first of these is human behaviour.

Often, the success of a project depends on the behaviour of certain people. For example, the aim of the project might be to reduce eutrophication in an urban river by having people reduce their use of fertilizers in home gardens, or to reduce air pollution by having factories install systems to remove pollutants from chimney emissions.

The issue is that, typically, not everybody cooperates with these sorts of projects. The degree of compliance varies from project to project, and this needs to be accounted for when we rank projects. Otherwise we risk giving funds to projects that have great potential but little benefit in practice.

Later on I'll discuss the estimation of adoption/compliance for particular projects. First I want to talk about how this information should be included in the project-ranking process.

To start with, define A as the level of adoption/compliance as a proportion of the level needed to achieve the project's goal. If $A = 0.5$, that means that compliance was only half the level we would have needed.

Usually, if A is less than 1.0, it doesn't mean the project generates no benefits. There is some relationship between A and the benefits generated. Figure 6 shows one possible example, where proportional benefits $[f(A)]$ increase slowly at low levels of adoption, then rapidly for a while, before flattening off again at high adoption. Other shapes are possible, but whatever the shape is, we know these important facts about it: it must range from zero (no adoption, so no project benefits) up to 1.0 (full adoption, so full project benefits). This follows from the fact that we define $f(A)$ as the proportion of target project benefits achieved.

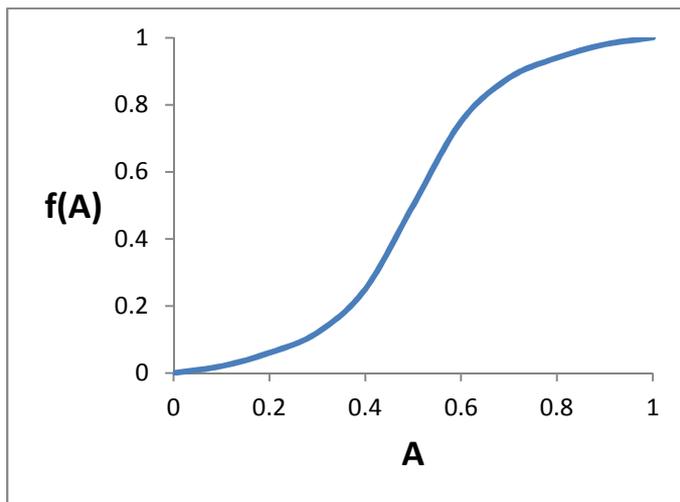


Figure 6.

This makes it obvious how $f(A)$ should be included in the formula we use for ranking projects: it should be multiplied by the potential benefits.

$$\textit{Benefit} = [V(P_1) - V(P_0)] \times f(A)$$

The terms in square brackets represent the difference in values between P_1 (physical condition with the project) and P_0 (physical condition without the project), assuming that there is full compliance, and this is scaled down by $f(A)$ to reflect the effect of less-than-full compliance. In other words, $[V(P_1) - V(P_0)]$ represents potential benefits, and we scale that down by $f(A)$ to get actual benefits – actual in the sense of accounting for lower adoption. (Equivalently, using the approach outlined in Section 3.3, $\textit{Benefit} = V(P') \times W \times f(A)$.)

[Note that if a project does not require anybody to change their behaviour, you would set $f(A) = 1$.]

This formula demonstrates an important principle for the ranking formula: if the benefits are proportional to a variable (as they are for $f(A)$), then that variable must be multiplied by the rest of the equation for benefits. Only that way can the formula correctly represent the reality that, if the variable is zero, the benefits must be zero, and if the variable is at its maximum value, so too are the benefits. As a way of testing whether this is relevant, ask this question: if the variable was zero, would the overall benefits be zero? If the answer is yes, the variable should probably be multiplied.

Unfortunately, a common way that people combine variables like these in the ranking formula is to give them weights (meant to reflect their relative importance) and add them up, something like this:

$$\textit{Benefit} = z_1 \times [V(P_1) - V(P_0)] + z_2 \times f(A)$$

where z_1 and z_2 are the weights. This is a bad mistake. With this formula, it is impossible to specify any set of weights that will make it represent the reality that the benefits are proportional to $f(A)$. Experiments I've done with this formula show that it can result in wildly inaccurate project rankings, leading to a big loss of environmental values. Switching from this bad formula to the correct one can be like doubling the program budget (in terms of the extra benefits generated) (see Pannell Discussion 158).

On the other hand, a simplification that is probably reasonable is to approximate $f(A)$ by a straight line. In practice, we usually have too little information about the actual shape of $f(A)$ in specific cases to be able to argue that its shape should be non-linear, and even if it is, it's unlikely to be so non-linear that an assumption of linearity would have very bad consequences. If you are comfortable with this approximation, you can just use A in the formula rather than $f(A)$.

$$\textit{Benefit} = [V(P_1) - V(P_0)] \times A$$

or

$$\textit{Benefit} = V(P') \times W \times A$$

That's what we usually do in INFFER; in the absence of better information, and in the interests of simplicity, we use A rather than $f(A)$ in the formula. But if somebody did have accurate numbers for $f(A)$, we would use them instead.

Finally, some brief comments on predicting the level of compliance/adoption for a project. There has been a great deal of research into the factors that influence the uptake of new practices (e.g., Rogers 2003; Pannell et al., 2006, and see www.RuralPracticeChange.org), so we have a good understanding of this. There are many different influential factors, and the set of important factors varies substantially from case to case.

However, despite the wealth of research, it remains difficult to make quantitative predictions about compliance for a specific project. One generalisation I would make is that people who develop projects are usually too optimistic about the level and speed of adoption that is realistic to expect – sometimes far too optimistic.

Specific predictions require specific knowledge about the population of potential adopters, and the practice we would like them to adopt. As far as I'm aware, there is only one tool that has been developed to help make quantitative predictions about adoption. This is ADOPT, the Adoption and Diffusion Outcome Prediction Tool. ADOPT is designed for predicting adoption of new practices by farmers. It is not suitable for other contexts, although it might provide insights and understanding that help people to make the required judgments.

3.5 Project risks

Environmental projects don't always go smoothly. There are various things that can stop them from delivering their intended benefits. One is the failure of enough people to change their behaviour in the desired ways, as discussed in Section 3.4. In this section, I'll discuss a number of additional risks that may affect projects. They need to be accounted for when ranking because projects vary greatly in how risky they are.

The word "risk" is used in so many different ways that it's important to be clear about what I mean by it. The risks I'm talking about here are things that might stop the project from succeeding, not the risks to the environmental assets (the threats that are making, or might make, the environmental assets degrade). I'll make some comments about those latter risks at the end of the section.

There are various types of project risks, potentially including:

- Technical risk (R_t): the probability that the project will fail to deliver outcomes for technical reasons. Management actions are implemented but they don't work because something breaks, or newly planted vegetation dies, or there was a miscalculation when designing the actions, or there is some sort of natural event that makes the actions ineffective.
- Social/political risk (R_s): the probability that social or political factors will prevent project success. For example, a project might rely on another government agency to enforce existing environmental regulations, but that agency is not prepared to enforce them because of the likelihood of a political controversy. Or there might be community protest, or perhaps even legal action, to stop the project.
- Financial risk (R_f): the probability that essential funding from partner organisations, or long-term funding for maintenance of benefits, will not be forthcoming. The latter one is often neglected. Many projects require ongoing funding for physical maintenance, or for continuing education or enforcement, without which the benefits would be lost. Often the decision to provide this ongoing funding is made independently of the decision to fund an initial project, so it is risky from the perspective of the funders of the initial project.
- Management risk (R_m): if different projects will be managed by different organisations, then there are likely to be differences in the risk of failure related to management. These risks might include poor governance arrangements, poor relationships with partners, poor capacity of staff in the organisation, poor specification of milestones and timelines, or poor project leadership.

All four of these risks can be important and are worth accounting for.

Some of these risks relate to all-or-nothing outcomes (e.g. there either is successful legal action against the project or there isn't), while others relate to continuous variables (e.g. maintenance funding might be deficient but not zero, resulting in some reduced level of ongoing benefits).

Representing risks for continuous variables is possible, but it requires fairly detailed information. Given that we are making educated guesses when we specify these risks, going to that level of detail is probably not warranted. What I suggest is to approximate each of the risks as the probability of a binary (all-or-nothing) variable turning out badly. To illustrate, rather than trying to specify probabilities for each possible level of maintenance funding, we would just specify the probability of maintenance funding being so low that most of the benefits would be lost. We would assume that there are two possible outcomes for each risk: it causes the project to fail, or the project is fully successful (or, at least, as successful as the other factors allow it to be).

Some risks might be correlated. For example, if there is social or political resistance to a project, it might reduce the probability of it getting long-term maintenance funding. In theory we should account for this correlation too, but again my view is that it is not worth going to that level of detail. Reasons include that: the quality of information we have when specifying these risks is not high; the formula used for ranking projects would have to get pretty complicated; and it would be confusing to many people.

Given those simplifications, the expected benefits of a project are proportional to the probability of the project NOT failing (1 minus the risk), for each of the separate risks. Again, proportional means multiplying, so:

$$\text{Expected benefit} = [V(P_1) - V(P_0)] \times A \times (1 - R_t) \times (1 - R_s) \times (1 - R_f) \times (1 - R_m)$$

or

$$\text{Expected benefit} = V(P') \times W \times A \times (1 - R_t) \times (1 - R_s) \times (1 - R_f) \times (1 - R_m)$$

With these variables included, the benefits are now probability-weighted, so they are "expected" benefits, in the statistical sense of a weighted average, where the "weights" are the probabilities of success (1 minus the probability of failure).

If you wanted to further simplify the approach, you could potentially combine all four of the risks into a single risk variable (R) representing the joint probability of the project failing for any of the four reasons.

$$\text{Expected benefit} = [V(P_1) - V(P_0)] \times A \times (1 - R)$$

That has the advantage of simplicity. Its disadvantage is that the individual risks tend to get a bit lost, and perhaps under-estimated, in the combined risk variable. In my view, it's worth taking the time to think separately about each of the risks, and if you do, you may as well have a variable for each.

Some organisations like to break risks down into likelihoods and consequences (as suggested in the ISO 31000 Risk Management Standard). Likelihoods represent the probability that a bad thing will happen (often scored on a scale like this: almost certain, likely, possible, unlikely or very unlikely), and consequence means how bad the bad thing would be if it did happen (e.g., scored as

insignificant, minor, moderate, major or catastrophic). Depending on the combination of these two scores, the overall risk is assessed as minimal, low, medium, high or extreme.

This is a rather drawn-out way of getting a risk score (compared with just stating the probability of project failure), and I don't think it's necessary, but it is logical and may help people to think clearly about the issues.

If you take that approach, a key question is, how should the overall risk score (minimal through to extreme) be used in the project ranking process? My recommendation is that you convert it into a probability of project failure. For example, you might specify that minimal risk corresponds to a 0.05 probability of project failure, low is 0.1, medium is 0.2, high is 0.5 and extreme is 0.8.

Having done that, you should use the probability in the equations I've given above. What you definitely should not do is give the risk number a weighting and add it onto (or subtract it from) the rest of the equation, but I've seen that done! Doing that implies that the losses from a poor outcome for a tiny project are just the same as for an enormous project, which is obviously wrong.

Finally, I want to return to a different usage of the word "risk", to mean a threat to the environmental asset. Environmental organisations sometimes conduct "risk assessments" in which they try to quantify the likely future extent of degradation from particular causes. In this document, we already dealt with that aspect of risk in Section 3.1; it represents the difference between current environmental condition and future condition without the proposed project ((1) – (3) in Figure 1).

A concern with this sort of risk assessment is that it may distract attention away from the correct measure of project benefits. Having done a "risk assessment" and come up with estimates of (1) – (3), people seem to find it difficult not to include them in the ranking formula. However, the correct measure of project benefits is (2) – (3), and if you include (2) – (3), also including (1) – (3) can only make the rankings worse. The point is that this type of "risk assessment" only provides the "without" half of the information you need to estimate potential project benefits. You also need to know what would happen "with" the project.

3.6 Time and time lags

Different projects involve different time lags until benefits are generated. There are at least four potential causes of lags:

1. Some projects take a significant amount of time to implement (implementation lags). For example, if a project relies on new research being conducted, it could take several years before results are available. Typically, implementation lags for different types of project range from a year to a decade.
2. It may take a while for the physical actions implemented in a project to take effect and to start generating benefits (effect lags). For example, if the project involves planting trees, they take a while to grow. Realistic effect lags in environmental projects range from zero to decades.
3. The project may be addressing a threat which has not occurred yet but is expected to occur in future (threat lag). For example, between 2001 and 2008 the Australian Government invested in many projects that were intended to prevent the occurrence of dryland salinity

in rural areas. In some of those areas, salinity was not predicted to occur until several decades in the future. In other cases, the threat lag was zero – the problem was already occurring.

4. If a project requires other people to change their behaviour or management, it may take a while for most people to change (adoption lag). Realistic adoption lags for substantial changes range from around five years (in exceptional cases) to several decades.

Given this variety of lags types, and the range of lag lengths within each lag type, projects vary widely in the overall time lag until benefits are generated. This makes it an important factor to consider when ranking projects but it's one that is commonly ignored.

Looking ahead from the time when an environmental project is being considered (time zero), a typical pattern of benefits over time (combining all the types of lags) is shown in Figure 7. This project generates half of its benefits by year 18 and 90% by year 25.

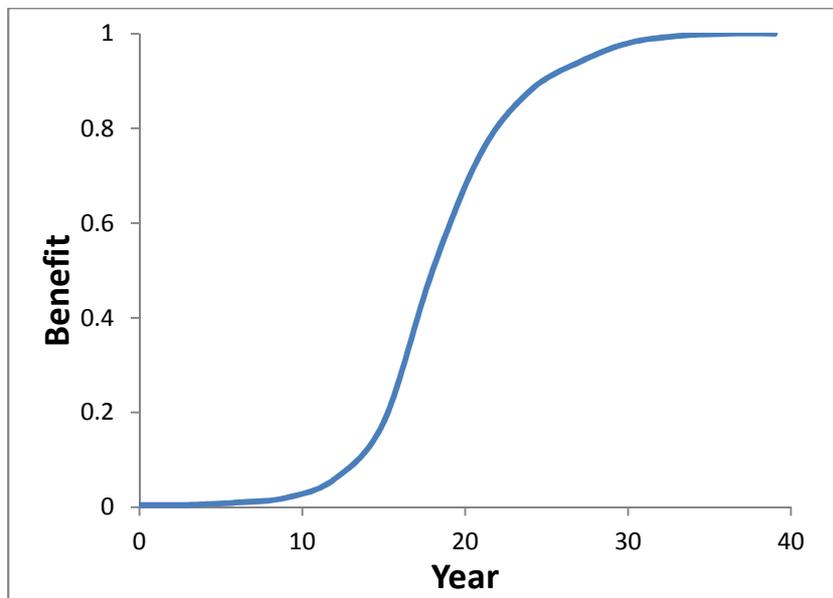


Figure 7.

The question is, how should differences in the overall time lags for different projects influence their rankings? Here's a simple example to illustrate how to think about it. (Inflation has been factored out of all the numbers in this example.)

Suppose there are two projects to rank. Both of them would involve costs of \$1 million in 2014 and both generate benefits worth \$2 million in the future. In one project, the \$2m benefits would be delivered in 2018 while in the other, the benefits would occur in 2033. The \$1m in 2014 has to be borrowed (at an interest rate of 5%) and will be repaid in full, including interest, when the benefits are generated. Thus, for the quick project, in 2018 we face a benefit (the \$2m) and a cost equal to the repayment. Similarly, for the slow project, in 2033 there is a benefit of \$2m and a different repayment cost. How do the costs and benefits stack up in each case?

The repayment cost would be calculated as follows:

$$\text{Future repayment} = \text{Present cost} \times (1 + r)^L$$

where *Present cost* is the amount borrowed up front, *r* is the interest rate (assumed to be 5%, after taking out inflation) and *L* is the time lag until the loan has to be repaid.

Compounding the interest costs over four years, the total repayment in 2018 would be \$1.22 m. This is less than the \$2 m benefit, so the quick project has a positive net benefit of \$0.78 m in 2018. On the other hand, by 2033, the repayment would be \$2.57 m, so the cost of the slow project would be greater than the benefit generated. Clearly, the quicker project would be preferred.

The same logic applies even if the money used to pay for the project doesn't have to be borrowed. The money used has an 'opportunity cost' (you miss out on the benefits of investing it in some other way) and that cost compounds over time in the same way as the interest on a loan.

The way that economists usually apply this thinking is to use the present as the reference date. Instead of compounding present costs into the future, they discount future benefits (and costs) back to the present. It amounts to exactly the same thing when it comes to ranking projects, and it has the advantage that it is easy to compare discounted values to the current value of money. The formula for present values is just the reverse of the repayment formula:

$$\text{Present value} = \text{Future value} / (1 + r)^L$$

Some people object to the idea of discounting future benefits, arguing that it is, in some sense, unfair or unreasonable. What they don't realise is that it's not really about the benefits – it's about the costs. Interest costs (or other opportunity costs) are real costs and shouldn't be ignored, but that's what you would be doing if you refused to discount benefits.

It's true that discounting benefits and costs in the distant future (e.g. 100 years) is more complicated, as issues of high uncertainty and inter-generational equity become important (e.g. see Pannell Discussion #34), but for shorter time frames (up to say 30 years) the logic behind discounting is robust (Pannell Discussion 224). A 5% real discount rate (with inflation factored out) is a pretty good general-purpose discount rate that's suitable for many environmental projects.

In principle, if we knew the year-by-year pattern of benefits (like in Figure 7), we would discount the benefit for each year and add them up to get the total present value of benefits. That would be the measure of benefits we used in the top line of the Benefit: Cost Ratio. If you have the required information, this is simple to do. Of course, getting the required information might not be simple.

In practice, there is usually a lot of uncertainty about how the benefits will play out over time. Recognising that uncertainty, it is probably usually not worth being too precise about the shape of the curve. A highly simplified curve, like the one in Figure 8, might be sufficient. All you need to know to specify this curve is the peak level of benefits (corresponding to the plateau in Figure 7), and the year when it will be achieved (corresponding to the year in Figure 7 when most of the benefits would be achieved).

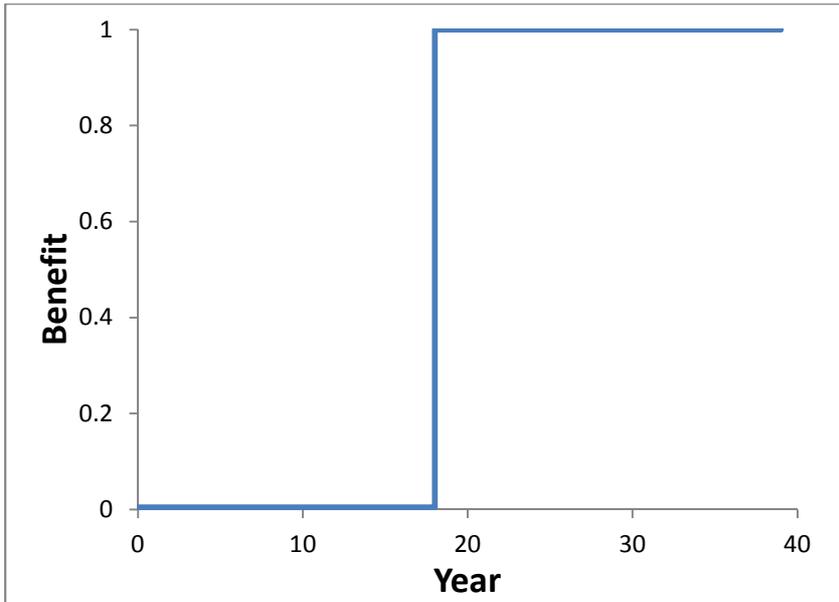


Figure 8.

This benefit curve is based on another convenient assumption: that the benefits, once, generated, will last forever. This is not too unreasonable if allowance is made for long-term maintenance funding (see Section 3.5 on project risks and a later section on project costs).

The simplified benefit curve has an advantage when it comes to calculating the present value of benefits. Rather than having to discount the benefits separately for each year (which is required if using Figure 7), we can just discount one number: the change in environmental values resulting from the project.

$$Benefit = [V(P_1) - V(P_0)] / (1+r)^L$$

This works because the overall value of the environmental asset ($V(P)$) consists of the discounted sum of its future benefits. So, looking at Figure 8, the change in value at year 18 consists of the discounted sum of benefits in all subsequent years. To convert that into a present value, we just need to discount them for another 18 years, which is what the above equation does if we set $L = 18$.

As discussed in previous sections, we also need to allow for adoption/compliance and project risks, so the equation for expected benefit is:

$$Expected\ benefit = [V(P_1) - V(P_0)] \times A \times (1 - R) / (1+r)^L$$

or

$$Expected\ benefit = V(P') \times W \times A \times (1 - R) / (1+r)^L$$

[I have combined the various project risks into one variable, R , to stop the equation getting too ugly. See Section 3.5.]

The time lag until benefits, L (= 18 in Figure 8) links back to the earlier section about measuring benefits as the difference in environmental value with and without the project (Section 3.1). In that section I noted that,

A practical simplification is to estimate the environmental benefits based on the difference in the asset value with and without the project in a particular future year. For example, we might choose to focus on 25 years in the future, and estimate values at that date with and without the project.

The selection of L tells us which particular future year to use for this with-vs-without comparison. So, for the example in Figure 8, the difference in values with and without the project would be estimated for year 18.

In Section 3.3 I discussed several possible ways to estimate and represent environmental values, some of which didn't involve expressing values in dollars (or Euros or Pounds or whatever). It might seem that discounting is not relevant to non-dollar values, but it is. Remember that discounting is used to account for the fact that an up-front cost grows over time due to compounded interest costs (or other opportunity costs), so it is applicable to any logical and consistent quantitative method for expressing future benefits. When comparing future benefits that occur at different times, you need to account for interest accumulating on up-front costs, even if the future benefits are not expressed in dollars. So you need to discount.

3.7. Multiple benefits

Previous sections have assumed that there is only one type of benefit being generated, or that multiple benefits generated by a project have already been converted into a common currency, such as dollars, and added up. What if we have multiple benefits and we want to account for them individually?

Suppose that a single project will generate three types of benefits, due to improvements to three different, but connected, environmental assets (e.g. a threatened species, a river that is suffering from reduced water quality, and an area of riparian vegetation that is attractive and provides habitat). The three assets have independent values V_1 , V_2 and V_3 .

All of the variables we've talked about in this document (the effectiveness of works, adoption, time lags and the various risks) could potentially have different values for each of the environmental benefits generated by this one project. For example, if the project is implemented, the risk of failing to improve water quality might be higher than the risk of failing to improve the condition of the riparian vegetation.

If the differences were significant enough, we might think it would be worth estimating three different values for each of the variables: W , A , L and the various R s.

If the V s were each measured in dollars, the expected benefits for the project would simply be the sum of the formulas for each benefit, as follows. [Three lots of the formula from Section 3.6.]

$$\text{Expected benefit} = V_1(P') \times W_1 \times A_1 \times (1 - R_1) / (1+r)_1^L +$$

$$V_2(P') \times W_2 \times A_2 \times (1 - R_2) / (1+r)^L_2 +$$

$$V_3(P') \times W_3 \times A_3 \times (1 - R_3) / (1+r)^L_3$$

To keep the formula simple, I've assumed that the relationship between adoption and benefits is linear, and I've combined the various risks into one overall probability of failure. Also, from now on I'm mostly going to use the $V(P') \times W$ version of the formula, rather than the $[V(P_1) - V(P_0)]$ version, which is equivalent. This has advantages in shaping the thinking, as outlined in Section 3.3, but also advantages for simplifying the formula where there are multiple benefits, as you'll see later.

If the values are not measured in money terms, you'll need to provide weights (z_1 , z_2 and z_3) to indicate the relative importance of the different benefits. The formula becomes:

$$\text{Expected benefit} = z_1 \times V_1(P') \times W_1 \times A_1 \times (1 - R_1) / (1+r)^L_1 +$$

$$z_2 \times V_2(P') \times W_2 \times A_2 \times (1 - R_2) / (1+r)^L_2 +$$

$$z_3 \times V_3(P') \times W_3 \times A_3 \times (1 - R_3) / (1+r)^L_3$$

This formula is getting pretty big and ugly. It also implies the need for a lot of information: the full equation for each type of benefit. Based on my experience, I'd say that most managers of real-world programs would not be prepared to go to this much detail. In reality, what commonly happens is that some of the variables are assumed to be the same for the different types of benefits. Often, I think that's a reasonable approximation of reality, or at least one that's not so bad that it's worth fighting against. If it seems reasonable to assume that W , A , R and L are the same for all three benefit types, then we can simplify the equation for expected benefits, as follows.

$$\text{Expected benefit} = [z_1 \times V_1(P') + z_2 \times V_2(P') + z_3 \times V_3(P')] \times W \times A \times (1 - R) / (1+r)^L$$

This is just the same as the formula in Section 3.6, but with $V(P')$ replaced by the large term in square brackets. Rather than being a single value, it's now the weighted sum of several values.

In previous sections in this document, I've been critical of the common practice of weighting and adding up variables in certain cases. However, this formula shows that it is not always a mistake. If we don't have dollar values, it's reasonable to weight and add the separate values to get an indicator of the total value at stake, prior to adjusting it down for W , A , R and L , as shown. The big mistake that is commonly made is to also weight and add W , A , R and L into the equation, rather than including them in the way shown above. Weighting and adding can be appropriate, but needs to be applied in a way that makes logical sense, rather than indiscriminately to all variables.

If we were weighting the $[V(P_1) - V(P_0)]$ version of the formula, it would look like this:

$$\text{Expected benefit} = \{z_1 \times [V_1(P_1) - V_1(P_0)] + z_2 \times [V_2(P_1) - V_2(P_0)] + z_3 \times [V_3(P_1) - V_3(P_0)]\} \times A \times (1 - R) / (1+r)^L$$

You can see that the other version is more compact.

Choices about the weights need to consider the way that the different benefits are scored. If the values are in dollars, all the weights become 1.0, so you end up just adding up the values.

If the values are not in money terms, the weights reflect the relative importance of the different benefits (a very subjective judgement), but they also need to account for the ranges over which the values are scored. For example, if value scores range from zero to 1.0 for one benefit but zero to 100 for another, the second one should probably have a much smaller weight to avoid it dominating the rankings. If the two benefits were equally important, the weight for the first one would need to be 100 times larger than the weight for the second one.

3.8 Private benefits

Suppose that an environmental project in a government program requires some private citizens to change their behaviour or business management. An example would be a project that aims to encourage farmers to convert crop land into native forest. As well as generating environmental benefits and public costs (the cost of running the project), this project may generate benefits and costs to the private citizens whose behaviour would need to change. I'll call these "private" benefits and costs.

Environmental projects are generally not set up with the aim of generating private benefits, but some do so as a spinoff from the real aim of generating environmental benefits for the broader public. How should the private benefits be accounted for?

The inclusion of private benefits is somewhat tricky and, in my observation, often poorly understood.

There are three reasons why the analyst might need to pay attention to private benefits and costs.

- (a) The level of private benefits influences the behaviour of the private citizens who the project aims to influence.
- (b) The level of private benefits should influence the choice of mechanisms used to try to encourage behaviour change.
- (c) It may be appropriate to account for private benefits as one of the benefits of the project, which may affect the ranking of projects.

(a) is pretty obvious, but surprisingly it often gets overlooked. When judging the level of adoption/compliance that is likely (Section 3.4), it's important to think about the practices being promoted and consider how attractive they are to the people they are being promoted to. If the private benefits of the new practice are much greater than the private costs, adoption will be high, and vice versa. It's not sufficient to consider the private benefits alone – you have to think about how they stack up against the private costs.

I should clarify that I don't just mean private financial benefits and costs here. Non-financial benefits and costs should be considered as well, potentially including factors such as riskiness, convenience, complexity, compatibility with existing practices, social pressures, and environmental outcomes that matter to the private person whose behaviour is potentially changing.

(b) is less obvious, perhaps, but also important. Agri-environmental programs in Australia have tended to over-rely on extension (information provision, awareness raising, etc.) to promote practices for which the private benefits are too low to outweigh the private costs. For those sorts of practices, something stronger, such as payments or regulation, would be needed to generate significant adoption, but of course these are more expensive approaches. For a project where private costs exceed private benefits, we face a trade-off that should be accounted for when the project is being evaluated. Extension is cheaper but generates little adoption, while payments or regulation generate more adoption but are more expensive.

In Europe, there is a tendency to use payments to farmers as the only response, even though extension might be sufficient (and much cheaper) in some cases. My Public: Private Benefits Framework is designed to help people think through these issues (Pannell, 2008).

(c) comes into play in some situations but not others. Including private benefits as a benefit of an environmental project is a can of worms, and they are unlikely to be big in most cases anyway, so the brief, pragmatic version of my advice is, leave them out, unless one of the objectives of the program is to generate private benefits. The next few paragraphs explain why I say that.

One reason you might want to leave them out is that they are not the main reason for doing the project. You might consider that the focus should be on maximising environmental benefits and that chasing private benefits when selecting projects might compromise that. Given that your organisation is likely to have a certain focus or set of responsibilities related to the environment, perhaps this is reasonable (provided that you don't also neglect private benefits when thinking about behaviour or mechanism choice).

However, economists generally prefer to consider all benefits and costs to the community as a whole. Most would consider that, if there are private benefits in excess of private costs resulting from a project, the net private benefits (private benefits minus private costs) should be counted as a benefit of the project.

But this is not as simple as it sounds. If the private benefits do exceed the private costs by a significant amount, one would expect the practice to be adopted by people anyway, even without the practice being promoted by the environmental program. In that case, there are no private benefits generated by the project. Indeed there would be no environmental benefits attributable to the project either, because the benefits would have occurred even without the project (the observed benefits are not "additional"). This gets back to being clear about what would happen with and without the project (Section 3.1).

An exception to this could be if the project makes the adoption occur sooner, by raising awareness of the new practice. The people who adopt sooner would benefit by getting the private benefits sooner than they would have (making them less affected by discounting – Section 3.6). While that is worth something, research I was involved in in agriculture suggests that the difference in adoption timing is likely to be only a year or two, in which case the additional private benefit would be quite small.

Some environmental programs provide payments to people to encourage them to change their behaviour. If the private costs of a new practice outweigh the private benefits, these payments can

help by offsetting losses that would have stopped participation. Should these payments be considered a private benefit and included in the project benefits formula? No, for the following reasons.

If the payments are pitched at just the right level to offset the private losses, then they exactly cancel out. There is no private net benefit to include. If the payments are bigger than they need to be, such that the recipients make a net benefit even after allowing for their private costs of participating, then in principle that net benefit could be included in the formula, but the excess payment will also have to be included as a cost to the project (we'll get to costs in Section 3.9), in a sense cancelling out the benefit. This additional cost is unnecessary and will reduce the funding available for other projects, so rather than worrying about accounting for these private benefits, the best response is to get rid of any private net benefits by making sure that any private payments are not excessive.

So, as I said earlier, the inclusion of private benefits as one of the benefits from an environmental project is complex, and the net private benefits are likely to be small – close to zero in many cases. Unless they are considered particularly important for some reason, or are one of the objectives of the program, it is probably reasonable to round them off to zero and leave them out of the analysis.

That's as much as I want to say about which variables should be included in the benefits part of the equation. I haven't discussed factors like whether the project fits within the program scope. I recommend applying that sort of factor as a filter on which projects get considered, rather than as a variable within the benefits formula.

3.9 Scoring variables

Let's return to the benefits part of the equation for ranking environmental projects (the simple version with only one benefit).

$$\text{Expected benefit} = [V(P_1) - V(P_0)] \times A \times (1 - R) / (1+r)^L$$

or

$$\text{Expected benefit} = V(P') \times W \times A \times (1 - R) / (1+r)^L$$

where V is the value of the asset (in whatever system of measurement makes sense), W is the effectiveness of works, A is adoption, R is risk and L is the time lag until benefits, in years. (For the purpose of this discussion, ignore r . It doesn't vary between projects. See Section 3.6.)

There are two distinct groups of variables in these equations. There are two variables that can take any value greater than zero: V and L . And there are three that can take any value between zero and one: W , A and R .

All of them are "continuous" variables – they change smoothly and can take any value within their feasible ranges. If you had the information, you would plug their exact values into the equation.

However, you never have exact information. A common approach in systems that collect information for ranking projects is to present a discrete number of options for the value of the variable, and ask participants to select the value that seems to be nearest to the correct value. For example, here is a question of this type about technical risk.

What is the probability that the benefits generated by the project would fall well short of expectations due to technical factors? (R_t)

- 0-5% Very low risk of project failure due to poor technical feasibility. ($R_t = 0.03$)
- 6-10% ($R_t = 0.08$)
- 11-15% ($R_t = 0.13$)
- 16-20% ($R_t = 0.18$)
- 21-100% High risk of long-term project failure due to poor technical feasibility. ($R_t = 0.60$)

I don't have a problem with this approach, as long as the response options are chosen thoughtfully. The quality of information available is usually not so high that this sort of approximation causes any significant reduction in the quality of the resulting rankings.

In the above example, I haven't spaced out the response options for R_t equally between zero and one, because I judged that most projects have values between zero and 20%. I have indicated the mid-point of the range for each response option, and that is what I would plug into R_t in the benefits equation.

Sometimes people convert the responses from this type of question into a number from an *ad hoc* scoring system, rather than using a scale that is more natural for the variable. For example, in the above case they might assign a score of 1 to the first response, 2 to the second response, and so on (instead of probabilities of 0.03, 0.08 and so on). This can potentially be OK, but there are a few traps to avoid.

Firstly, in a case like the one above where the response options are not equally spaced, the scores assigned should not be equally spaced either. They should be spaced out consistent with the values in the response options.

Secondly, if one of the response options represents zero, the score assigned to that option should be zero. For example, if a response option for adoption is zero adoption, it should get a score of zero so that when it is multiplied into the equation, the overall score is zero. (Obviously, if there is zero adoption of the actions being promoted by the project, there would be no environmental benefits attributable to the project.) In that case, using scores of 1, 2, 3, 4 or 5 is no good. If you must use scores instead of probabilities, use 0, 1, 2, 3 or 4 (assuming that the first response option is zero adoption).

Thirdly, even if you are using an *ad hoc* scoring system, you still have to multiply the variables, as shown in the above equations. Weighting them and adding them up, as done in many systems, produces much inferior rankings.

If benefits aren't being measured in dollars, then a scoring system can work, as long as it satisfies the above requirements. However, my advice is to use the correct ranges when scoring each variable (i.e. between zero and one for W , A and R) rather than some *ad hoc* system. Why not do that? It is no more difficult, it makes it easy to meet all the above requirements, and it makes the meaning of each variable clearer.

If the benefits are being measured in dollars, then you don't have an option. You have to assign values that correspond to the meanings of the variables rather than using an *ad hoc* scoring system. Otherwise you lose the benefit of being able to assess whether the benefits exceed the costs.

4. Costs

Remarkably, some systems for prioritising environmental projects only look at the benefits and ignore the costs. In certain special cases, this can be OK, but in most cases it's a serious mistake. Ironically, although it might feel like focusing on the environmental benefits should be good for the environment, neglecting costs results in a (usually large) reduction in the total environmental benefits generated from a given pool of funding.

Even where costs are included in the ranking process, people tend to think that the cost aspect is simple and straightforward. In fact there are some tricky aspects to be aware of.

There are several different types of costs related to environmental projects that may need to be considered when ranking projects. They are: the cost of the project itself (cash costs to the funder, in-kind costs to the lead organisation and private costs to participants), ongoing costs to maintain the benefits generated by the project, and the opportunity cost of the funds invested in the projects. The last one is different in nature from the others, and we have already dealt with it in Section 3.6 (on time lags and discounting). First, I'll cover project cash costs.

4.1 Cash costs

For an environmental project to proceed, money is needed for various purposes, potentially including: the time of people employed to implement the project or provide support to the project, cars, fuel, machinery and equipment, payments to people to encourage behaviour change, legal costs, office space, telephones, insurance, publicity, printing, and so on.

Whether an organisation is allocating its own cash resources amongst projects or applying to an external funder, the cash allocated to a project may not be enough to pay for all associated costs. Often the lead organisation bears some of the costs out of existing salary or operating budgets (in-kind costs), or people in the community voluntarily contribute time or other resources (private costs). I'll cover these non-cash costs in the next two sections.

Accounting for new cash allocated to a project is straightforward. It counts as a cost, and should be included in the denominator of the Benefit: Cost Ratio.

$$BCR = \frac{V(P') \times W \times A \times (1 - R)/(1 + r)^L}{C}$$

where C is total project cash costs.

In theory, to calculate C , costs in future years should be discounted to their present values (see Section 3.6) before adding them up. How important it is to do this discounting depends on the number of years over which the funding will be rolled out. If it is only a short project (e.g., three years), the error in total costs if you ignore discounting of costs is only a few percent, so it doesn't matter much. Leaving cash costs undiscounted might even help to offset the common tendency for excessive optimism about projects (see Pannell Discussion 213).

On the other hand, if the project funding is longer in duration, discounting probably matters. For example, if the project involves constant funding over 10 years, failing to discount would exaggerate total costs by around 25% (assuming a 5% discount rate). For some projects, this would affect whether benefits exceed costs (which can only be seen if benefits are being measured in dollar terms).

Discounting costs matters most if some of the projects being considered involve funding over much longer time frames than others, because in those cases it can affect the ranking of projects.

Cash costs unambiguously go in the denominator, because the level of cash available for funding environmental actions within a program is always limited. Whenever projects require a resource that is limited in availability, the right ranking is calculated by dividing project benefits by the quantity of that resource required by each project.

I've seen a ranking metric in actual use that put a weight on cash costs and subtracted them from the estimate of project benefits. This is slightly better than ignoring costs entirely, but only slightly. There's no reason not to get this right – it's trivially easy.

4.2 In-kind costs

If the cash funds provided for a project are not sufficient, it's common for the organisation that's running the project to make up the difference. For example, they may cover some or all of the costs of: project staff salaries, administrative support, office space, stationary, telephone calls, and so on.

These 'in-kind' costs provided by the organisation responsible for the project are real costs and should be accounted for when projects are being ranked.

Generally, organisations have a limited level of resources that they can use to provide in-kind support to projects. From a theoretical perspective, this creates a slight problem, because there are now two limited pools from which different types of costs (cash and in-kind) are being drawn.

Strictly speaking, this means that you cannot rank projects based on a single formula. The BCR only ranks projects correctly if there is only one constraint on costs (e.g. on the level of funds in the pool of new cash). If there are two constraints, it is impossible to specify a formula that is guaranteed to rank projects correctly. In theory, what one would have to do is build a constrained optimisation model with separate constraints for the two types of funding (e.g. using the technique of mathematical programming).

In practice, that is almost never going to happen, because of the additional time and expertise required, and the loss of transparency and intuitiveness in the results. Fortunately, treating all costs as if they come from one large pool (i.e. including them all in the denominator) is usually a very good approximation of the theoretically correct approach. (The loss of environmental benefits was less than 1% in the examples I looked at.) So that's the approach I recommend – use ...

$$BCR = \frac{V(P') \times W \times A \times (1 - R)/(1 + r)^L}{C + K}$$

where K is in-kind costs provided by the organisation.

As was the case with cash costs, if in-kind costs are borne for more than a few years, then they should be discounted before adding them up to calculate K (Section 3.6).

There is one combination of circumstances when it would be OK to ignore in-kind costs. This is when (a) in-kind costs are close to being the same proportion of cash costs for every project, and (b) you only care about the ranking of projects, not in assessing whether their overall benefits exceed their overall costs. Situation (b) is reasonably common, but I don't think you would often find situation (a), especially if projects from different organisations are to be compared when funding decisions are made.

4.3 Private costs

People or businesses who participate in environmental projects may bear costs. For example, people may contribute their time, their land, or other inputs to the project.

The question of whether and how to consider these costs when ranking projects is not as straightforward as for the other costs we've discussed.

First consider the case where private costs are contributed voluntarily. Just as I argued for leaving out private benefits (Section 3.8), my advice is to leave out these voluntary private costs.

The reason is that, since people are doing it voluntarily, if they actually do it, it must be something they want to do. For them, personally, the private benefits (broadly defined) must be sufficient to outweigh the private costs. In other words, overall, there are no net private costs to include. If there were private net costs from participation (i.e. if private costs exceeded private benefits), then they would not participate, and no private costs would be incurred.

[Just to be clear, in those comments about the balance between private costs and private benefits, I don't just mean private financial benefits and costs, but also all of the psychological and social benefits and costs that people get from contributing to an environmental project.]

Things are different, however, when the private costs are not borne voluntarily, but are imposed on people against their will, such as through enforcement of a regulation. In that case, it's quite possible that private costs exceed private benefits and we have "compliance costs" that should be accounted for. My simple advice is to estimate them, discount them if necessary, and add them in to the denominator of the BCR, like this:

$$BCR = \frac{V(P') \times W \times A \times (1 - R)/(1 + r)^L}{C + K + E}$$

where E represents total compliance costs.

The compliance costs you need to estimate are the losses incurred by the relevant people, aggregated over the whole group of people, and aggregated over the relevant time frame (including discounting if the time frame is more than a few years).

Compliance costs may continue on longer than the initial project costs if people are required to continue doing something that they wouldn't otherwise choose to do. Over what time frame should they be counted? See the next section on maintenance costs.

Compliance costs should be calculated as the change in private net benefits with and without the project. They should include any direct and indirect losses, and any offsetting benefits, if relevant. Estimating compliance costs can be quite challenging unless you have good survey information or an economic model to use, but it's better to make an educated guess than to leave them out.

The advice to add compliance costs to the denominator is another case of weighing up the benefits of practical simplicity with the costs of an approximation. If you're not a stickler or an enthusiast for theoretical detail, you might not want to bother with the rest of this section where I present my justification for this simple advice. If you are interested, read on.

As I've said before, putting costs in the denominator of the formula used to rank projects is the right approach when there is a limited pool of resources to be allocated amongst projects. The limited pool might be held within the organisation delivering the projects, or in a separate funding organisation.

On the other hand, if supply of a costly input is not in limited supply, then in principle its costs does not belong in the denominator. Rather, the cost should be subtracted from the numerator.

Given that, what should we do with compliance costs? Should they be added to the denominator or subtracted from the numerator? If there is effectively a constraint on the total level of compliance costs the community will bear, it is reasonable to add them to the denominator. (This approximates the effect of optimising for an additional constraint, just as it did for the organisation's in-kind costs.)

If there is no such constraint, then in theory they should be subtracted from the numerator. But what if the numerator is not measured in dollars? I suppose you could apply a weight to the compliance costs (with another subjectively determined weight) and subtract the weighted cost, but I feel that's likely to be even more of an approximation than adding them to the denominator, so I'd do the latter.

That leaves us with compliance costs being subtracted from the numerator only in the case where benefits are measured in dollars and there is no constraint on the level of compliance costs. However, if you're dealing with a variety of different projects, including compliance costs differently in different cases might be more confusing than it's worth. If we were to simplify the system by always adding compliance costs to the denominator, the error would usually be small, in my

judgement. That's why this is my "simple advice" on how to include compliance costs: add them to the denominator.

4.4 Maintenance costs

Often, environmental projects need ongoing funding in the long term to preserve or maintain the benefits generated by an initial project. For example, funds may be needed to maintain, repair, or replace equipment or structures; to pay the wages of people responsible for ongoing education, training or enforcement; or for continuing payments to people to ensure ongoing adoption of improved environmental practices. These costs might arise for a few years beyond the end of the initial project, or they might last more-or-less forever.

The required level of maintenance funding can be substantial, potentially exceeding the cost of the initial project, so it's an important factor that needs to be accounted for when ranking projects. However, it usually isn't! I've never seen any system for ranking environmental projects that does account for it, other than ones I've helped develop.

How should it be included? First, define M as the level of maintenance funding that would be required to fully maintain the project's benefits in the long term. Because maintenance costs tend to be required for a long time, they need to be discounted before they are added up. If M_3 , for example, is the maintenance cost in year 3, then the total discounted maintenance cost is given by ...

$$M = M_1/(1+r)^1 + M_2/(1+r)^2 + M_3/(1+r)^3 + \dots$$

where r is the real discount rate.

If maintenance costs have to be continued into the indefinite future, the question arises, how long should the time frame be for calculating M ? There is no clear-cut answer to this. The length of time used for the calculations needs to be fairly long, but if it's extremely long, discounting means that maintenance costs in the distant future are quite insignificant in the present. Also, uncertainty about what might happen in the distant future is very high, so one might judge that it's not worth factoring in benefits or costs that may never arise in reality. My suggestion is to use a time frame of around 25 years, although I couldn't argue strongly against a somewhat shorter or significantly longer time frame.

So with discounted maintenance costs included, the formula for the BCR becomes ...

$$BCR = \frac{V(P') \times W \times A \times (1 - R)/(1 + r)^L}{C + K + E + M}$$

There is one final refinement to make to the BCR formula. I included several risks in the benefits part of the equation (Section 3.5), summarised into one risk, R , in the above equation. Of the four risks I included, one of them may also have an impact on the cost part of the equation. This is R_f , the probability that required maintenance funding will not be forthcoming. "Required" in this context means that most project benefits would be lost in the absence of maintenance funding.

Failing to receive maintenance costs has impacts on two of the cost variables, M and probably E . It's obvious that if no maintenance costs are received, M would be zero. Therefore, we should weight M by the probability that maintenance costs will be received, $(1 - R_f)$.

Compliance costs might occur during the initial project phase, or during the maintenance phase, or both. The component that would occur in the maintenance phase should also be weighted by the probability that maintenance costs will be received, because if they aren't received, the project will presumably collapse, and there will be no enforcement of compliance. In the version of the equation below, for simplicity I've assumed that all compliance costs occur in the maintenance phase. The equation also includes, for the first time since Section 3.5, all of the risks shown separately (as in Pannell et al., 2013).

$$BCR = \frac{V(P') \times W \times A \times (1 - R_t) \times (1 - R_s) \times (1 - R_f) \times (1 - R_m) / (1 + r)^L}{C + K + (E + M) \times (1 - R_f)}$$

[R_f also includes the probability that a partner organisation will not deliver essential resources that it agreed to provide, resulting in project failure. I'm assuming that a result of that would be that costs E and M would not be incurred.]

Up to now I've used M to represent the full required level of maintenance cost. What if some maintenance funding is likely, but it's expected to be insufficient to fully maintain project benefits in the long term? You might want to adjust three variables. Firstly, you would reduce M to the expected level of funding. Secondly, you might want to reduce R_f to reflect the fact that obtaining the lower level of maintenance funding is easier. And thirdly, the benefits should be scaled down to some degree (by reducing the estimate of W , or reducing $[V(P_1) - V(P_0)]$). How much they should be scaled down depend on how sensitive the benefits are to a reduction in maintenance funding.

For a good project, providing sufficient maintenance funding is likely to increase benefits by more than it increases costs. If it doesn't, then it indicates that the proposed investment in maintenance is excessive.

The equation above is the final new version I'll show in this paper (although there are several sections still to come). This version provides a comprehensive, logical, theoretically respectable, and practical equation for ranking projects. It embodies quite a few simplifications, but none that are likely to have more than a minor adverse impact on the ranking results. It avoids a number of other simplifications (and errors) that are likely to have very serious impacts on the rankings.

I'll come back and summarise the formula and its components and rationale in the last section of the document. Before then, there are a couple of more issues related to costs to cover, and then some high-level issues to discuss, including uncertainty, the issue of using simplifications, and key mistakes to avoid.

4.5 Sharing costs

This issue relates to the sharing of costs between different benefits. In Section 3.7 I talked about how to assess projects that generate multiple benefits. A related issue I struck once was in a case where the organisation wanted to rank potential investments in a number of threatened species

individually, even though they knew that the actions needed to protect one species would help to protect others as well. I can see why they would want to do this – it would be tidy to be able to create a ranked list of all the species.

The approach I suggested to them was to define S as the share of total costs that is attributable to the current species, and base it on the share of benefits. You would add up all the benefits for different species resulting from the actions taken to protect this species, and then ask what share of the benefits belongs to this species? Then that share, which is S , gets multiplied by total costs in the BCR for this species.

Generally, I wouldn't recommend this approach unless it's important to create a ranked list of each individual environmental asset. For that purpose, it is probably the best that can be done, but it's still a somewhat crude approximation. It's better to rank projects rather than assets (see Section 1) and if a project generates multiple benefits, so be it – use one of the approaches in Section 3.7.

5. Uncertainty

Uncertainty and knowledge gaps are unavoidable realities when evaluating and ranking environmental projects. The available information is almost always inadequate for confident decision making. Key information gaps often include: the cause-and-effect relationship between management actions and environmental outcomes; the likely behavioural responses of people to the project; and the environmental values resulting from the project – what is important or valuable about the environmental outcomes and how important or valuable are they?

It has been argued to me that uncertainty about the data and objectives is generally so high that it is not worth worrying too much about the procedure used to prioritise projects. Any procedure will do. If that was really true, no analysis could help with decision making – we might as well just draw projects out of a hat.

In fact, while it's true that uncertainty is usually high, it's not true that the ranking procedure doesn't matter, particularly when you consider the outcomes across a portfolio of projects. Even given uncertain data, the overall environmental benefits of a program can be improved substantially by a better decision process. Indeed, environmental benefits appear to be more sensitive to the decision process than to the uncertainty. For example, I have found that there is almost no benefit in reducing data uncertainty if the improved data are used in a poor decision process (Pannell 2009). On the other hand, even if data is uncertain, there are worthwhile benefits to be had from improving the decision process.

This is certainly not to say that uncertainty should be ignored. Once the decision process is fixed up, uncertainty can make an important difference to the delivery of environmental benefits.

There are economic techniques to give negative weight to uncertainty when ranking projects. I've used them and I think they are great for research purposes. However, I don't recommend them for practical project-ranking systems. They aren't simple to do properly, so they add cost and potentially confusion.

Instead of representing uncertainty explicitly in the ranking equation, I suggest a simpler and more intuitive approach: rating the level of uncertainty for each project; and considering those ratings subjectively when ranking projects (along with information about the Benefit: Cost Ratio, and other relevant considerations).

Apart from its effect on project rankings, another aspect of uncertainty is the question of what, if anything, the organisation should do to reduce it. In my view, it is good for project managers to be explicit about the uncertainty they face, and what they plan to do about it (even if the plan is to do nothing). Simple and practical steps could be to: record significant knowledge gaps; identify the knowledge gaps that matter most through sensitivity analysis (Pannell, 1997); and have an explicit strategy for responding to key knowledge gaps as part of the project, potentially including new research or analysis.

In practice, there is a tendency for environmental decision makers to ignore uncertainty when ranking projects, and to proceed on the basis of best-guess information, even if the best is really poor. In support of that approach, it is often argued that we should not allow lack of knowledge to hold up environmental action, because delays may result in damage that is costly or impossible to reverse. That's reasonable up to a point, but in my view we are often too cavalier about proceeding with projects when we really have little knowledge of whether they are worthwhile. It may be at the expense of other projects in which we have much more confidence, even though they currently appear to have lower BCRs. It's not just a question of proceeding with a project or not proceeding – it's a question of which project to proceed with, considering the uncertainty, environmental benefits and costs for each project. When you realise this, the argument based on not letting uncertainty stand in the way of action is rather diminished.

In some cases, a sensible strategy is to start with a detailed feasibility study or a pilot study, with the intention of learning information that will help with subsequent decision making about whether a full-scale project is worthwhile, and how a full-scale project can best be designed and implemented. A related idea is "active adaptive management", which involves learning from experience in a directed and systematic way. Implementation efforts get under way, but they are done in a way which is focused on learning.

Particularly for larger projects, my strong view is that one of these approaches should be used. I believe that they have great potential to increase the environmental benefits that are generated. They imply that the initial ranking process should not produce decisions that are set in stone. Decisions may need to be altered once more information is collected. We should be prepared to abandon projects if it turns out that they are not as good as we initially thought, rather than throwing good money after bad.

As far as I'm aware, the sorts of strategies I'm suggesting here are almost never used in real-world environmental programs. Managers are never explicit about the uncertainties they face, there usually isn't a plan for addressing uncertainty, projects are funded despite profound ignorance about crucial aspects of them, proper feasibility assessments are never done, active adaptive management is almost never used, and ineffective projects that have been started are almost never curtailed so that resources can be redirected to better ones. In these respects, the environment sector is dramatically different from the business world, where people seem to be much more concerned about whether their investments will actually achieve the desired outcomes. Perhaps the difference

is partly because businesses are spending their own money and stand to be the direct beneficiaries if the investment is successful. Perhaps it's partly about the nature of public policy and politics. Whatever the reason is, I think there is an enormous missed opportunity here to improve environmental outcomes, even without any increase in funding.

6. Simplifications

Throughout this document, I've struck a balance between theoretical accuracy and simplifications to make the process more practical and less costly. Clearly, this balance involves judgement. Others might judge that more or fewer simplifications are appropriate, or prefer different simplifications than the ones I've recommended. One thing they would have to agree on, though, is that simplifications are essential to make the system workable. The ones I've recommended are carefully chosen, on the basis that they are unlikely to have serious impacts on the total level of environmental benefits generated by a portfolio of projects. In some cases, the careful choosing I've done is based not just on subjective judgement, but on numerical analysis.

Even with the simplifications I've suggested, the process is still rather information hungry. If dealing with a large number of potential projects, collecting all the information for all of the projects may be more costly than is warranted, especially if the level of funding available is small relative to the total cost of all potential projects. For example, I've worked with environmental bodies which had upwards of 500 potential projects to rank, but were likely to get funding for less than 5 per cent of them.

In this type of situation, it is justifiable to use an even more highly simplified approach initially to filter down the full list of 500 projects to a manageable number for more detailed (but still simplified) assessment. An approach I've found effective is to select a few of the most important variables (e.g. the importance or significance of the environmental assets affected; the likely technical effectiveness of management actions; the likely degree of cooperation with the project by those people or businesses whose behaviour would need to change). Each project is scored for each of these key variables on a simple three- or four-level scale (low, medium, high, or very high). Then one looks for projects with three scores of high or better. If that doesn't provide a sufficient number of potential projects, loosen the criterion a bit: look for projects with two scores of high or better and one medium. Loosen or tighten as needed to get a workable number of projects to assess further. Projects that meet the criterion you end up settling on go through more detailed assessment using the BCR equation and the rest of the projects are put aside.

Clearly, with such a simplified process, there is a chance that good projects will be rejected or poor projects will be let through. As long as enough good projects get through the initial filter, missing some good ones is not likely to be a big problem. And as long as the projects that pass through the filter are subjected to the more detailed assessment, letting poor projects through is not a problem at all (apart from wasting some time) because they will be rejected following the detailed analysis.

Now let's come back to the simplifications included in the detailed BCR calculation. Most of them have been spelled out in previous sections. Key simplifications that I judge to be reasonable in most cases include:

- Assuming that environmental benefits are linearly related to the proportion of people who adopt the desired new practices or behaviours;
- Representing project risks as binary variables: success or complete failure;
- Having only one time lag for all benefits from the project;
- Approximating the private benefits and voluntary private costs as zero; and
- Treating the project costs, maintenance costs and compliance costs as if there was only one combined constraint on their availability

There are also a few other simplifications that I haven't mentioned so far, but which are implicit in the equations I've presented in earlier sections. I've had the first two of these pointed out by economists with eyes for theoretical detail, and the third by a colleague with a particular interest in this issue.

Firstly, I've been assuming that the value of an environmental asset does not depend on the conditions of other related assets. In reality, the benefits of project A could depend of whether project B is funded and, if so, there is no definitive ranking of individual projects. In practice, the error resulting from my simplifying assumption is likely to be small enough to ignore. Pretty much everybody who ranks environmental projects makes this assumption, and ignores any error. But if the issue is judged to be important enough to be worth accounting for, you could define a project that combines the activities of projects A and B into one project and compare its BCR with those for project A and project B individually.

Secondly, if one assumes that projects are defined at a particular scale and cannot be scaled up or down, then ranking using the BCR may not be accurate because it doesn't account for the risk of leaving some of the funds unspent. [This is known as the "knapsack problem".] That's true, but unless funding is sufficient for only a small number of projects, the loss from ranking using the BCR is likely to be very small. For example, Hajkowicz et al. (2007) estimated losses of between 0.3% and 3% in a particular program. And if you abandon the normally unrealistic assumption that the scale of each project is fixed, then the losses disappear almost entirely. When you factor in the transaction costs of building, solving and explaining a mathematical programming model to solve a knapsack problem properly, you would always rank by BCR.

Thirdly, the equations I've presented only measure benefits arising directly from the project. Graham Marshall pointed out that participation in a current project might also generate benefits for future projects by building mutual trust and networks amongst the participants (i.e., "social capital"). He even experimented with simple ways to estimate this benefit so that it could be added to the equation. Unfortunately, the feedback from participants in Graham's experiments was that accounting for this benefit added significantly to the complexity of the process. Furthermore, my judgement is that, while these are real benefits, they are probably not usually large enough or different enough between projects to make a notable difference to project rankings. For that combination of reasons, I haven't included them.

7. Mistakes to avoid

Prior sections in this document have mostly focused on things that should be done when ranking environmental projects. Now and then I've commented on things that should not be done, but this time that is the main focus. The mistakes I describe here are all things that I've seen done in real systems for ranking projects.

Weighting and adding. If you've read the whole thing, you are probably sick of me saying not to weight and add variables, except in particular circumstances (Section 3.7). I'm saying it one more time because it is such a common mistake, and one with such terrible consequences. I've had someone argue that despite all the logic, weighting and adding should be done for all variables because it gives decision makers scope to influence the results to reflect their preferences and values, thereby giving them ownership of the results. Absolute nonsense. That's like giving people the flexibility to make up their own version of probability theory. There is no benefit in them owning the results if the results are really bad. There are much better ways to give influence to decision makers, such as by allowing them to adjust the value scores (V) to reflect their judgements about what is important. Doing it by weighting and adding together the wrong variables introduces huge errors into the results and greatly reduces the environmental values generated by a program.

Including "value for money" as a criterion separate from the variables that determine value for money. This seems to be quite common too. A number of times I've seen systems that ask questions about relevant variables (like environmental threats, adoption, values, risk, costs) but then have a separate question about value for money, rather than calculating value for money based on the other information that has already been collected. This is unfortunate. A subjective, off-the-top-of-the-head judgement about value for money is bound to be much less accurate than calculating it from the relevant variables. This behaviour seems to reveal a lack of insight into what value for money really means. If the aim is to maximise the value of environmental outcomes achieved (as it should be), then value for money is the ultimate criterion into which all the other variables feed. It's not just one of the criteria; it's the overarching criterion that pulls everything else together to maximise environmental outcomes.

Here's a recent experience to illustrate what can go wrong. I was asked to advise an organisation about their equation for ranking projects. They had specified the following as separate criteria for selecting projects: value for money, logical consistency of the project, and likelihood of successful delivery of the project. But, of course, the logical consistency of the project, and the likelihood of successful delivery are both things that would influence the expected value for money from the project. They are not distinct from value for money, they are part of it. I would consider them when specifying the level of risk to include in the equation. Specifically, they determine the level of management risk, R_m (Section 3.5).

Unfortunately, somebody in the organisation who had power but no understanding insisted that logical consistency and successful delivery be treated as criteria at the same level as value for money, and worse still that they all be weighted and added! My explanations and protests were dismissed. As a result, they lost control of their ranking formula. Rankings for small projects were determined almost entirely by the scores given for logical consistency and successful delivery, and barely at all by the Benefit: Cost Ratio (BCR), and the rankings for large projects were the opposite – completely unaffected by logical consistency and successful delivery. The ultimate result was poor project rankings, leading to poor environmental outcomes.

Messing up the with-versus-without comparison. Back in Section 3.1 I talked about how the benefits of a project should be measured as the difference in outcomes between a world where the project is implemented and a world where it isn't ($[V(P_1) - V(P_0)]$ or W). When you say it like that, it sounds like common sense, so it's surprising how many systems for ranking projects don't get this

right. Some don't include any sort of measure of the difference that a project would make. They may use measures representing the importance of the environmental assets, the seriousness of the environmental threats, or the likely level of cooperation from the community, but nothing about the difference in environmental values resulting from the project.

Some systems include a difference, but the wrong difference. I've seen a system where the project benefit was estimated as the difference between current asset condition and the predicted asset condition if nothing was done (current versus without). And another which used the difference between current asset condition and predicted asset condition with the project (current versus with). Both wrong.

Finally, I've seen a system which did include the correct with-versus-without difference, but still managed to mess it up by also including a couple of inappropriate variables: current asset condition, and the current-versus-without difference. In this situation, more information is not better – it will make the rankings worse.

Omitting key benefits variables. Because the benefits part of the equation is multiplicative, if you miss out one or more of its variables, the inaccuracies that are introduced are likely to be large. If you ignore, say, adoption, and projects vary widely in their levels of adoption, of course it's going to mean that you make poor decisions.

Ignoring some or all of the costs. Almost all systems ignore maintenance costs. Most ignore compliance costs. Some ignore all costs. Some include costs but don't divide by them. All mistakes.

Failing to discount future benefits and costs. Another very common mistake – a variation on the theme of ignoring costs.

Measuring activity instead of outcomes. If asked, pretty much everybody involved in ranking environmental projects would say that they want the resources they allocate to achieve the best environmental outcomes. So it's frustrating to see how often projects are evaluated and ranked on the basis of activity rather than outcomes. For example, benefits are sometimes measured on the basis of the number of participants in a project. This ignores critical factors like the asset values, the effectiveness of the on-ground works, and the project risk. Sometimes this approach arises from a judgement that participation has benefits other than the direct achievement of outcomes. No doubt, this is true to some extent. In particular, participation by community members in a current project can build "social capital" that reduces the cost of achieving environmental outcomes in subsequent projects. In Section 6 I recorded my judgement that measuring that particular benefit is probably not worth the trouble in most cases (at least for the purpose of ranking projects). The reasons are that it's a somewhat complex thing to measure, and that those indirect benefits would usually not be large enough or different enough between projects to affect project rankings much. I'm making a judgement here, of course, but I think it is irrefutable that considering only activity/participation and failing to estimate direct benefits due to improved environmental outcomes is likely to compromise project rankings very seriously. But that does sometimes happen.

Negative scores. This is a really strange one that I don't expect to see again, but I mention it because it was a catalyst for writing this paper. I was once involved in a project ranking process where the organisation was scoring things using an *ad hoc* points system. Most variables were being scored on

a five-point scale: 1 for the worst response through to 5 for the best. The designers of the process decided that they'd penalise projects that were rated "high" or "very high" for risk by extending the range of scores downwards: -5 (for very high risk) to +5 (for very low risk). They were using the dreaded weighted additive formula and, naturally enough, the weighting assigned to risk was relatively high, reflecting their view of its importance. This was in addition to risk having the widest range of scores. They didn't realise that combining these approaches would greatly amplify the influence of risk, with the result that project rankings depended hugely on risk and not much on anything else. At the meeting, someone from the organisation commented that risk was dominating the ranking, but they couldn't understand why. Others agreed. I explained what was going on and advised them that their system would have been more transparent and easier to control if they had left the range of scores the same for each variable and just varied the relative weights.

That experience highlighted to me how very little some people who design ranking systems understand about what they are doing. This document is an attempt to provide an accessible and understandable resource so that if people want to do a good job of the ranking process, they can. In the next section I'll provide a summary of the whole argument.

8. Summary

Around the world, thousands of different quantitative systems have been used to rank environmental projects for funding. It seems that every environmental body creates anew or re-uses at least several such systems each year. Judging from the examples I have examined, most of the systems in use are very poor. The performance of many of them is not much better than choosing projects at random. If only people would be more logical and thorough in their approach to ranking environmental projects! The potential to reduce wastage and improve environmental outcomes is enormous. That's why I wrote this paper.

There are many ways that you can go wrong when putting together a formula to rank projects, and unfortunately the quality of the results is quite sensitive to some of the common errors. Common important mistakes include: weighting and adding variables that should be multiplied; messing up the comparison of outcomes with versus without the project; omitting key benefits variables; ignoring costs; and measuring activity instead of environmental outcomes.

Fortunately, though, it's not hard to do a pretty good job of project ranking. A bit of theory, some simple logic and a dose of common sense and judgment lead to a set of specific guidelines that are presented in this document. The essential points are as follows.

1. The core criterion for ranking projects is value for money: a measure of project benefits divided by project-related costs. This is the criterion into which all the variables feed. It's how you pull everything together to maximise environmental outcomes.
2. You should rank specific projects, rather than environmental assets. You cannot specify numbers in the ranking formula for some of the key variables without having in mind the particular interventions that will be used.
3. There are always many different ways of managing an environmental asset, and they can vary greatly in value for money. Therefore, it can be worth evaluating more than one project per asset, especially for large, important environmental assets.

4. Benefits of a project should be estimated as a difference: with versus without the project, not before versus after the project.
5. Weak thinking about the "without" scenario for environmental projects is a common failing, sometimes leading to exaggerated estimates of the benefits.
6. There are two parts to a project's potential benefits: a change in the physical condition of the environment, and a resulting change in the values generated by the environment (in other words, the value of the change in environmental services).
7. Those potential benefits usually need to be scaled down to reflect: (a) less than 100% cooperation or compliance by private citizens or other organisations; (b) a variety of project risks; and (c) the time lag between implementing the project and benefits being generated, combined with the cumulative cost of interest on up-front costs (i.e. "discounting" to bring future benefits back to the present).
8. If in doubt, multiply. That's a way of saying that benefits tend to be proportional to the variables we've talked about (or to one minus risk), and the way to reflect this in the formula is to multiply by the variables, rather than weighting and adding them. Don't take this too literally, however. You can mess up by multiplying inappropriately too.
9. Weighting and adding is relevant only to the values part of the benefits equation (when there are multiple benefits from a project), not to any other part.
10. Don't include private benefits as a benefit or voluntary private costs as a cost, but do include involuntary private costs as a cost.
11. Other costs to include are project cash costs, project in-kind costs, and maintenance costs (after the project is finished). Costs get added up, rather than multiplied.
12. Uncertainty about project benefits is usually high and should not be ignored. The degree of uncertainty about each project should be considered, at least qualitatively, when projects are being ranked. Also, decisions about projects should not be set in stone, but modified over time as experience and better information is accumulated. Strategies to reduce uncertainty over time should be built into projects (e.g. feasibility assessments, active adaptive management).
13. Where the cost of all projects that are in contention greatly exceeds the total budget, it is wise and cost-effective to run a simple initial filter over projects to select a smaller number for more detailed assessment. It's OK to eliminate some projects from contention based on a simple analysis provided that projects are not accepted for funding without being subjected to a more detailed analysis.

There are a number of simplifications in the above advice. Simplifications are essential to make the system workable, but care is needed when selecting which simplifications to use.

In summary, the content and structure of the ranking formula really matters. A lot. A logical and practical formula to use is:

$$BCR = \frac{V(P') \times W \times A \times (1 - R_t) \times (1 - R_s) \times (1 - R_f) \times (1 - R_m) / (1 + r)^L}{C + K + (E + M) \times (1 - R_f)}$$

where

BCR is the Benefit: Cost Ratio,

$V(P')$ is the value of the environmental asset at benchmark condition P' ,

W is the difference in values between P_1 (physical condition with the project) and P_0 (physical condition without the project) as a proportion of $V(P')$,

A is the level of adoption/compliance as a proportion of the level needed to achieve the project's goal,

R_t, R_s, R_f and R_m are the probabilities of the project failing due to technical risk, socio-political risks, financial risks and management risks, respectively,

L is the lag time in years until most benefits of the project are generated,

r is the annual discount rate,

C is the total project cash costs,

K is the total project in-kind costs,

E is total discounted compliance costs, and

M is total discounted maintenance costs.

V can be measured in dollars, or in some other unit that makes sense for the types of projects being ranked. The advantages of using dollars are that it allows you to (a) compare value for money for projects that address completely different types of environmental issues (e.g. river water quality versus threatened species) and (b) assess whether a project's overall expected benefits exceed its total costs.

For some projects, it works better to calculate potential benefits in a different way: $[V(P_1) - V(P_0)]$ rather than $V(P') \times W$. They are equivalent but involve different thought processes.

A simplification that might appeal is to combine all four risks into one overall risk, R . If you do that, also drop ' $\times (1 - R_f)$ ' from the denominator (because you no longer have a separate number for R_f). This simplification makes the formula look a bit less daunting, but it probably doesn't really save you any work, because you should still consider all four types of risk when coming up with values for R .

This formula works were there is a single type of benefit from a project, or where the V scores for multiple benefits have already been converted into a common currency, such as dollars, and added up. If a project has multiple benefits and you want to account for them individually, replace $V(P')$ by the weighted sum of the values for each benefit type. For example, if there are three types of benefits, use $[z_1 \times V_1(P') + z_2 \times V_2(P') + z_3 \times V_3(P')]$, where the z 's are the weights. I'm assuming here that the other benefit variables (W, A, R and L) are the same for each benefit type. If that's not approximately true, you need to adjust the formula further.

One reaction I get is that it all looks too complicated and surely isn't worth the bother. My response is to ask, if you could double your budget for projects by putting a bit more effort into your project ranking process, would you do so? Of course you would. Doubling the environmental benefits generated from your environmental investments is rather like doubling your budget. If your current

ranking system is of the usual questionable quality, doubling the benefits (or more) is readily achievable by using the approaches advocated here.

That's all! Thanks for reading and best of luck with your project ranking endeavours.

References

- Gauvin, F.-P. (2009). *What is a Deliberative Process?* National Collaborating Centre for Healthy Public Policy, Quebec, http://www.ncchpp.ca/docs/DeliberativeDoc1_EN_pdf.pdf
- Hajkowicz, S., Higgins, A., Williams, K., Faith, D.P. and Burton, M. (2007). Optimisation and the selection of conservation contracts, *Australian Journal of Agricultural and Resource Economics* 51(1), 39-56.
- Jenkins, G.P., Kuo, C.-Y. and Harberger, A.C. (2011). Discounting and alternative investment criteria, Chapter 4 in *Cost-Benefit Analysis For Investment Decisions*, IDEAS page for this paper. <http://ideas.repec.org/p/qed/dpaper/197.html>
- Office of Best Practice Regulation (2009). Best Practice Regulation Guidance Note, Decision rules in regulatory cost-benefit analysis, Australian Government, Department of Finance and Deregulation, <http://www.finance.gov.au/obpr/docs/Decision-Rules.pdf>
- Pannell, D.J. (1997). Sensitivity analysis of normative economic models: Theoretical framework and practical strategies. *Agricultural Economics* 16(2), 139-152.
- Pannell, D.J. (2008). Public benefits, private benefits, and policy intervention for land-use change for environmental benefits, *Land Economics* 84(2), 225-240.
- Pannell, D.J. (2009). The cost of errors in prioritising projects, INFFER Working Paper 0903, University of Western Australia. <http://dpannell.fnas.uwa.edu.au/dp0903.htm>
- Pannell, D.J. and Roberts, A.M. (2010). The National Action Plan for Salinity and Water Quality: A retrospective assessment, *Australian Journal of Agricultural and Resource Economics* 54(4), 437-456.
- Pannell, D.J., Marshall, G.R., Barr, N., Curtis, A., Vanclay, F. and Wilkinson, R. (2006). Understanding and promoting adoption of conservation practices by rural landholders. *Australian Journal of Experimental Agriculture* 46(11), 1407-1424.
- Pannell, D.J., Roberts, A.M., Park, G., Alexander, J., Curatolo, A. and Marsh, S. (2012). Integrated assessment of public investment in land-use change to protect environmental assets in Australia, *Land Use Policy* 29(2), 377-387.
- Pannell, D.J., Roberts, A.M., Park, G. and Alexander, J. (2013). Designing a practical and rigorous framework for comprehensive evaluation and prioritisation of environmental projects, *Wildlife Research* 40(2), 126-133.
- Roberts, A.M. Pannell, D.J. Doole, G. and Vigiak, O. (2012). Agricultural land management strategies to reduce phosphorus loads in the Gippsland Lakes, Australia, *Agricultural Systems* 106(1), 11-22.
- Rogers E.M. (2003). *Diffusion of Innovations*, 5th ed., Free Press: New York.